

REVIEW

Deconstructing the sources of genotype-phenotype associations in humans

Alexander I. Young^{1,2*}, Stefania Benonisdottir¹, Molly Przeworski^{3,4*}, Augustine Kong^{1*}

Efforts to link variation in the human genome to phenotypes have progressed at a tremendous pace in recent decades. Most human traits have been shown to be affected by a large number of genetic variants across the genome. To interpret these associations and to use them reliably—in particular for phenotypic prediction—a better understanding of the many sources of genotype-phenotype associations is necessary. We summarize the progress that has been made in this direction in humans, notably in decomposing direct and indirect genetic effects as well as population structure confounding. We discuss the natural next steps in data collection and methodology development, with a focus on what can be gained by analyzing genotype and phenotype data from close relatives.

Not long ago, genetic analyses were performed using trait values (phenotypes) in families, without genetic data. The discovery of readily measurable genomic markers enabled the identification of disease genes by linkage analysis, without prior knowledge of the underlying mechanisms (1). This led to the identification of the gene responsible for the X-linked phagocytic disorder chronic granulomatous disease in 1986, followed by those for other Mendelian diseases such as cystic fibrosis (2) and Huntington's disease (3) as well as the breast cancer genes (4, 5). This approach was also applied to the study of common complex diseases, including type 2 diabetes, but failed to provide replicable findings.

The second major development came with high-throughput single-nucleotide polymorphism (SNP) arrays, which allowed for the genotyping of hundreds of thousands of SNPs simultaneously, giving rise to the genome-wide association study (GWAS) (6). A GWAS tests each SNP for association with the phenotype, without family data. The success of GWAS started with the discovery that *CFH* contributes to age-related macular degeneration; that analysis was based on 96 cases and 50 controls (7). Subsequent increases in sample size, with some now more than 2 million (8), have led to the discovery of thousands of genetic variants affecting hundreds of human traits. Results from GWAS hold the promise of identifying novel drug targets (9, 10), among other applications.

The power of a GWAS to identify a trait-affecting SNP depends on the fraction of trait variation explained by the SNP, which increases

in proportion to the square of the effect size and the heterozygosity. Because heterozygosity is higher for more common variants, initial successes were mainly for susceptibility variants with a minor allele frequency above 5%. Even if a common variant is not directly analyzed, it is likely to be strongly correlated with a genotyped SNP nearby, because of the lack of ancestral recombination events between them. This correlation is termed “local” linkage disequilibrium (LD). Non-local LD—correlations between variants that are not physically close—can result from nonrandom mating. As a result of local LD, GWAS usually does

‘For many complex traits, GWAS has changed the landscape of genetic investigations and our understanding of genetic architectures’

not directly identify the specific causal variant, but only localizes its approximate genomic position. Fine-scale mapping, which often requires functional analysis and experimentation, is needed to identify causal variants (11).

The majority of common variants found by GWAS to affect disease risk have low to modest effects (increasing the odds of disease by less than a factor of 1.5 per risk allele) (12, 13). Application of GWAS to whole-exome and whole-genome sequencing, along with statistical imputation of sequence-level variants into samples genotyped by SNP arrays, has led to the discovery of some rarer variants with large effects (14). Although the trait variance explained by genome-wide significant (GWS) loci has increased, for most complex traits, the variance explained by GWS loci is only a fraction of the estimated heritability. This gap,

labeled the “missing heritability,” is discussed below (12, 15).

For complex traits, identifying all the causal variants and elucidating their underlying mechanisms remains a distant goal. However, GWAS data can be used for prediction from genotypes, notably with polygenic scores (PGS). PGS combines the estimated effects of multiple genetic variants to provide a predicted trait value for an individual. Many applications of PGS have been investigated, such as identification of individuals with substantially elevated genetic risk of heart disease (16). Despite the demonstrated value of PGS, questions regarding robustness and interpretation (i.e., what is driving the predictive power) have started to surface (17, 18).

In GWAS, it is widely acknowledged that associations can be biased by population stratification, primarily association between ancestry and environment effects. Methods adjusting for ancestry, together with replication (19), lend confidence that most GWS associations with common SNPs are true positives. However, this does not mean that bias is eliminated, nor the nature of genotype-phenotype associations properly characterized. We aim to lay out here the different contributions to genotype-phenotype association, explain the difficulties they introduce, and propose possible solutions.

Effects captured by GWAS associations

The association between a genetic variant and a phenotype can be decomposed into the direct effect of the variant, the indirect genetic effect of the variant, and confounding effects (Fig. 1). An example would be a variant that has a direct effect on educational attainment (EA) when inherited, as well as an indirect effect through parental behavior/nurture (20). The same variant could also have an indirect effect on health through parental nurture, but little to no direct effect. Direct effects incorporate a wide range of causal pathways, some neither simple nor “direct”; for example, variants in *CHRNA5* affect lung cancer risk through their association with smoking quantity (21). Furthermore, the direct effect here can include effects of other variants in local LD. Note that typical GWAS conducted without family data can only estimate the sum of the direct and indirect effects (combined effect), not the two separately.

Under an additive model for the joint effects of variants, we define a genetic component as a linear combination of the genotypes of all the causal variants with weights proportional to the true (direct, indirect, or combined) effects (Fig. 2). The genetic components for direct effect and indirect effect are distinct, but they can be correlated with a strength that depends on the genetic correlation between the proband phenotype of interest and the phenotypes of the relatives through which the indirect effects are mediated. As an example, this correlation is probably strong for EA and weak for body mass index (BMI) (20). The relative strengths of these two genetic components and their correlation determine the correlations with the combined-effect genetic

¹Big Data Institute, Li Ka Shing Centre for Health Information Discovery, University of Oxford, Oxford, UK. ²Social, Genetic and Developmental Psychiatry Centre, Institute of Psychiatry, Psychology and Neuroscience, King's College London, London, UK. ³Department of Biological Sciences, Columbia University, New York, NY, USA. ⁴Department of Systems Biology, Columbia University, New York, NY, USA.

*Corresponding author. Email: alextiyoung@gmail.com (A.I.Y.); mp3284@columbia.edu (M.P.); augustine.kong@bdi.ox.ac.uk (A.K.)

Spectrum of genetic ancestries among families (population structure)

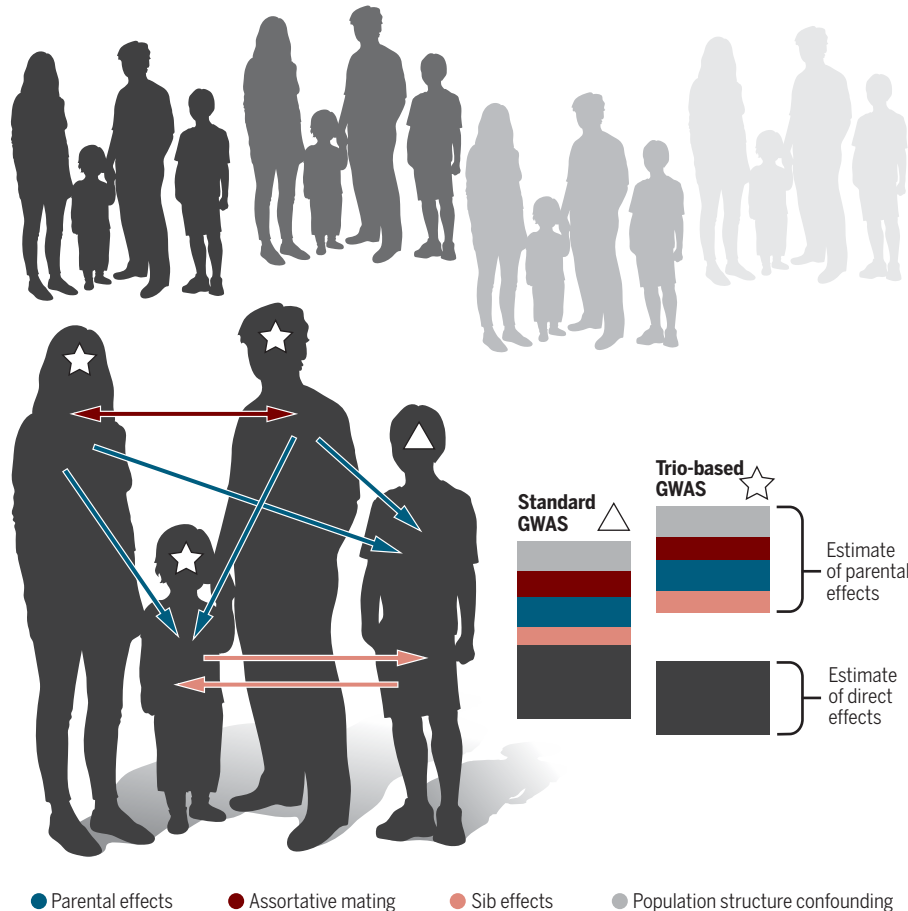


Fig. 1. The signals captured by GWAS of distantly related individuals and families. When based on distantly related individuals, estimates of effect sizes of SNPs on a trait include direct genetic effects (black) as well as a number of other effects, including confounding due to population structure (gray), assortative mating for the trait or a correlated one (burgundy), indirect genetic effects from parents (blue), and sib effects (peach). Family-based GWAS (such as the use of a trio) uses parental genotypes as controls to separate direct from indirect genetic effects and other confounding effects (20), as illustrated in the decomposition to the right. In this figure, we ignore effects of local LD.

component. Because the PGS constructed from a typical GWAS uses estimates of the combined effects, its predictive power can sometimes be substantially stronger than what can be explained by the direct effects alone (20).

Genetic effects can contribute to the associations between traits through pleiotropy. A two-trait model of pleiotropy (Fig. 2, top) of the combined effects has three parameters: the variances explained by the combined-effect genetic components of the two traits, and the correlation between them. This correlation has been estimated for many pairs of traits using GWAS data (22). By separating out direct- and indirect-effect components, the model (Fig. 2, bottom) has 10 parameters, including the magnitudes of four direct- and indirect-effect genetic components, and six correlations. The full model cannot be estimated using standard GWAS, so we currently have little understanding of the extent to which

direct and indirect genetic effects contribute differently to pleiotropy.

Confounding effects

The association between a genetic variant and a phenotype could reflect, in part, a correlation with some other causal phenomenon (environmental or genetic) rather than a true causal effect of the SNP on the phenotype. This type of confounding arises from the presence of non-random mating leading to population structure. There are at least three sources of confounding in GWAS: (i) environmental confounding, where allele frequencies and environmental effects vary in a correlated way across different geographic regions or subpopulations; (ii) genetic confounding, when allele frequency differences between subpopulations correlate with frequency differences of other alleles with causal effects; or (iii) assortative-mating confounding, which occurs

when there is assortative mating for the trait (or a correlated trait), a variant with a causal effect on the trait becomes correlated with other variants with causal effects, and its association with the trait then captures its own causal effect plus a fraction of that of the other variants. These forms of confounding are conceptually different, but in practice they are often intertwined.

Adjusting for confounding in GWAS

Principal component (PC) adjustment is a common technique used to remove some of the population structure-related confounding effects (23). Ideally, the principal components used for adjustments are strongly correlated with the environmental confounding component and uncorrelated with the direct genetic effect component. If the direct effect component is substantially correlated with the confounding components, PC adjustment will remove some of the direct genetic effects as well as confounding effects.

The assortative-mating confounding component (iii) is, by its nature, nearly perfectly correlated with the sum of the direct and indirect components. Assortative mating for traits such as height and EA (24) leads to nonlocal LD of variants with direct and indirect effects, which PCs capture. Thus, in theory, PC adjustment could adjust away most of the direct effect component. In practice, this does not happen. Even with a very large sample size, the inferred PCs are likely to be mostly noise beyond a few strong (often geographic) signals. Results from the UKB white British (WB) sample highlight this point (Fig. 3): Beyond the first eight strongest PCs, PCs computed from a sample of 272,519 individuals (25) appear to be mainly driven by sampling noise and local LD within chromosomes. The noise can mask subtle population structure that can lead to confounding in GWAS even after PC adjustment (26).

Fitting linear mixed models (LMMs) is an alternative to PC adjustment. These methods perform a type of regression on a set of SNPs where the effect of each SNP is modeled as a “random effect” drawn from a normal distribution (27). LMMs have long been used for trait prediction in animal breeding (28). In human studies, LMM association testing typically consists of estimating the effect of a focal SNP as a “fixed effect” while modeling random effects for a set of other SNPs. Naïve LMM computation scales with the cube of sample size, and thus alternative computational approaches have been developed to handle large GWAS sample sets (29).

The appeal of LLMs is that they enable improved modeling of population stratification and sample relatedness (27). LMMs are often used in combination with PCA adjustment and can account for more complicated patterns of stratification by modeling the effects of (nearly) all measured SNPs, capturing both real genetic effects and stratification effects (27). Furthermore, LMM methods can lead to improved estimation of SNP effects and their sampling errors over linear regression in the presence of sample relatedness (27). LMMs can also reduce bias in SNP effect estimates due to assortative mating (30).

Box 1. Glossary.

Assortative mating: When couples that produce offspring select one another on the basis of particular phenotypes.

Fine-scale mapping: Refers to approaches that aim to identify which variant or variants are likely to be causal among the set of associated variants identified in a GWAS.

Heritability: Measures the proportion of phenotypic variation explained by the direct effects of all genetic variants in a population at a given time.

Heterozygosity: The probability that two alleles at a site differ; assuming Hardy-Weinberg equilibrium and considering a biallelic site, this measure of genetic diversity is given by $2p(1-p)$, where p is the allele frequency.

Imputation: A statistical method that infers the genotypes of individuals at variants not directly measured on a genotyping array by reference to complete genome sequence data.

Indirect genetic effect: The effect of a genetic variant in one individual on the trait of another individual through the environment.

Genome-wide significant (GWS) associations: Variants associated with the phenotype at a significance level chosen to overcome the multiple testing burden, usually set at $P < 5 \times 10^{-8}$.

Linkage analysis: Tests for cosegregation of phenotypes and genotypes within families.

Pleiotropy: The common observation that many SNPs that are associated with one trait are also associated with other traits. Related to the concept of genetic correlation.

Principal component: A principal component is an inferred axis of genetic variation in a sample. A principal component is a linear combination of genotypes of SNPs, where each SNP has a “loading” giving its contribution to the principal component.

Polygenic score (PGS): Weighted sum of alleles carried by an individual, where the weights are given by effect sizes estimated in GWAS.

However, current LMM GWAS methods do not remove the contribution of indirect genetic effects.

Using family genotype data

Given the parental genotypes, an offspring's genotype is determined by random segregation of genetic material during meiosis. This random segregation is uncorrelated with indirect genetic effects from relatives and other confounding effects. Parental genotypes can thus be used as controls to obtain unbiased estimates of direct genetic effects (20, 31) (Fig. 1). Similarly, genetic differences between siblings are a result of random Mendelian segregation in the parents during meiosis. The genetic differences between siblings are therefore not confounded with indirect genetic effects from parents, population stratification, and assortative mating. However, methods using the differences in sibling genotypes estimate the direct effect minus the indirect effect from the sibling, and hence only provide unbiased direct-effect estimates when the indirect genetic effect of the sibling is zero.

The study of indirect genetic effects has a long history in animal breeding (13). In humans, most studies of indirect genetic effects have used PGS derived from GWAS that do not distinguish between direct and indirect genetic effects (Fig. 4) (20). However, when direct and indirect genetic effects are not perfectly correlated, that approach

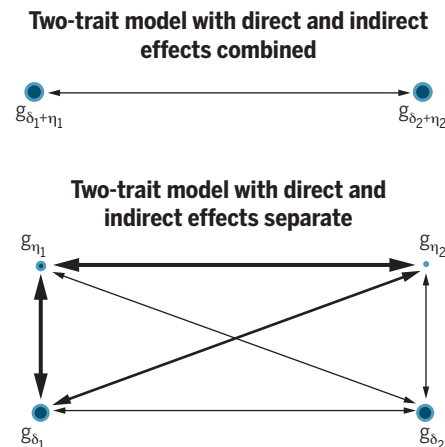


Fig. 2. Two-trait genetic models with direct and indirect effects combined or separated.

For a trait, assuming an additive model, the genetic component combining direct and indirect effects is $g_{\delta+\eta} = \sum_i (\delta_i + \eta_i) g_i$, where g_i , δ_i , and η_i denote the genotype, the direct effect, and the indirect effect of variant i , respectively. Top: With two traits (1 and 2), there are two magnitudes and one correlation. For each trait, the combined genetic component can be separated into the direct-effect component, $g_{\delta} = \sum_i \delta_i g_i$, and the indirect-effect component, $g_{\eta} = \sum_i \eta_i g_i$. Bottom: The two-trait model becomes one with four genetic components and six pairwise correlations between them. For the canonical example illustrated here, where trait 1 could be EA and trait 2 could be BMI, the size of a dot indicates the magnitude of a component, and the thickness of a connecting line indicates the strength of the correlation.

cannot give the full picture (32). Ideally, GWAS should be performed with parental or sibling genotypes as controls and using models with indirect genetic effects. However, the power of this approach is currently limited because large samples with genotyped siblings and/or parents are uncommon. Furthermore, as only around half of the genetic variation in a population is within-family, substantially larger samples of families are required to obtain the same study power as standard GWAS analysis. Therefore, methods combining information from standard GWAS and from analysis of families are needed.

Heritability

Traditionally, heritability has been estimated from comparing correlations between identical and non-identical twins. In addition to identifying specific causal loci, it is possible to use GWAS data to estimate the phenotypic variation explained by the genetic variation captured by the SNPs (and variants in LD with them) on a genotyping array, called “SNP heritability” or h_{SNP}^2 (33). Estimates of h_{SNP}^2 imply that common genetic variants assayed on a typical genotyping array collectively explain substantially more phenotypic variance than the GWS variants. However, estimates of h_{SNP}^2 tend to be substantially lower than estimates of heritability from twin studies (15), part of the “problem of missing heritability.” Some, but far from all, of this gap is explained by effects of imputed variants that are not in strong LD with markers on a typical genotyping array (13, 34). One possibility is that much of the remaining missing heritability is explained by very rare variants (35).

A widely used method, GREML, estimates h_{SNP}^2 by measuring the strength of the relationship between phenotypic similarity and genome-wide genetic similarity (estimated from SNPs), which varies even for the distantly related individuals typically used in GWAS (36). This approach provides an estimate of the total variance explained by the combined direct and indirect effects of probands' alleles (20, 31). The extent to which indirect genetic effects and population stratification have contributed to estimates of h_{SNP}^2 (Fig. 4) is not known, nor is the bias induced by assortative mating on both within- and between-family estimates of heritability.

It is also important to note that the total variance explained by the combined direct and indirect effects differs from the traditionally defined heritability, which is about direct effects only. However, it is a parameter of interest, as it defines an upper bound of genetic prediction from probands' alleles. An implication is that the upper limit of genetic prediction for a trait could often be larger than the heritability (18).

Some recent methodological developments**LD score regression**

With the explosion of GWAS, approaches have been developed to better use and interpret their results. Notably, LD score regression (LDSC) was developed to distinguish the effects of confounding

due to population stratification from causal genetic effects on GWAS test statistics (37). Assuming a highly polygenic architecture, the GWAS test statistic for an individual SNP is expected to increase with its LD score (a measure of the genetic variation tagged by a SNP through local LD) because of increasing correlation with causal variants. However, the average test statistic across all SNPs is raised by population stratification, due to correlation between alleles and differences in mean trait values between subpopulations (37–39). By estimating how much population stratification-induced confounding inflates the average test statistic, the LDSC intercept can be used to adjust the GWAS test statistics. LDSC can also be used to estimate the correlation between SNP effects on different traits (22), to partition contributions to SNP heritability from different functional categories of variants (40), and to facilitate multitrait meta-analysis (41).

A key assumption of LDSC is that allele frequency differences between subpopulations are independent of LD scores (37). However, a correlation between LD scores and allele frequency differences can be induced by forms of linked selection such as background selection (26). Thus, questions remain about the reliability of the LDSC measure of population stratification bias.

Mendelian randomization

Mendelian randomization (MR) uses genetic data to improve causal inference in epidemiology (42). If a genetic variant affects trait A, and trait A affects trait B, then variants that affect trait A are expected to affect trait B. Genetic variants that affect trait A can be used to determine whether an association between trait A and trait B reflects a causal influence of trait A on trait B, given that the genetic variants affect trait B only through their effect on trait A, and that the genetic variants are not correlated with any confounding factors. MR has proven successful in refuting false causal hypotheses derived from observational data, such as the association between HDL cholesterol levels and cardiovascular disease (43) and the reduced risk of cardiovascular disease in moderate drinkers in Western societies (44).

MR usually relies on SNP effect estimates from GWAS without families, which can be biased by population stratification, indirect genetic effects from relatives, and assortative mating (45). Within-family MR methods have been proposed to address these concerns and have shown that previous MR estimates of causal effects of height and BMI on EA were spurious (45).

A further challenge for MR analyses is widespread pleiotropy: If a SNP affects trait B through a trait other than trait A, then it is not a valid instrument for inference of the causal effect of trait A on trait B. Although methods have been developed to address this problem, their effectiveness can depend on prior knowledge about the confounding pathway (46).

Gene-by-environment interactions

A gene-environment (GxE) interaction occurs when a genetic variant's effect on a trait differs in

different environments (47). Such GxE interactions are distinct from gene-environment correlation, which can result from, for example, indirect genetic effects from relatives. In humans, robustly replicated examples of GxE interactions are rare outside of pharmacogenomics (48, 49). One exception is an interaction between variants in the *FTO* locus and physical activity affecting BMI (50, 51).

The power to detect GxE interactions in GWA studies is likely to have been low as a result of small effect sizes and multiple-testing burden. One way to increase the power to detect GxE is to look for interactions between environmental factors and PGS (52, 53). This method is effective when genetic variants affecting a trait interact with environmental factors in similar ways, but cannot identify interactions between environmental factors and specific genetic variants. LMMs can be applied to detect a component of phenotypic variance arising from the interaction between genome-wide genetic variants and an environmental factor (54) but cannot pinpoint interactions with specific genetic variants. Genetic variants involved in GxE interactions affect the variability of a trait (55, 56), which can be exploited to reduce the search space of potential interactions by restricting to variants with evidence for an effect on phenotypic variability. However, methodological challenges remain: Interaction effects and genetic effects on phenotypic variability are

sensitive to the scale of measurement (56, 57), and the effects of population stratification on estimates of GxE are not well characterized. Furthermore, the causality of GxE interaction effects is hard to establish, because the interaction may be with an unmeasured environmental factor that is correlated with measured environmental factor(s), and the broader socioenvironmental factors that may structure the environmental exposure are often unknown.

Portability of phenotypic prediction

The accuracy of prediction based on PGS depends on the trait's heritability and the power of the existing GWAS (notably on the sample size and genetic architecture) (28). For a handful of traits [such as height, for which the current prediction accuracy is ~25% (58)], existing scores are already informative in sets of individuals similar to those in which the GWAS was conducted.

Polygenic scores do not perform as well in predicting phenotypes of individuals that differ from those included in the GWAS set. Some of the reasons are understood and arise from differences in ancestry. Notably, because PGS consists of a weighted sum of allele counts and because allele frequencies vary across the globe (due to genetic drift and natural selection), alleles that contribute to trait variation in the GWAS are less likely to be present or may even be absent in more

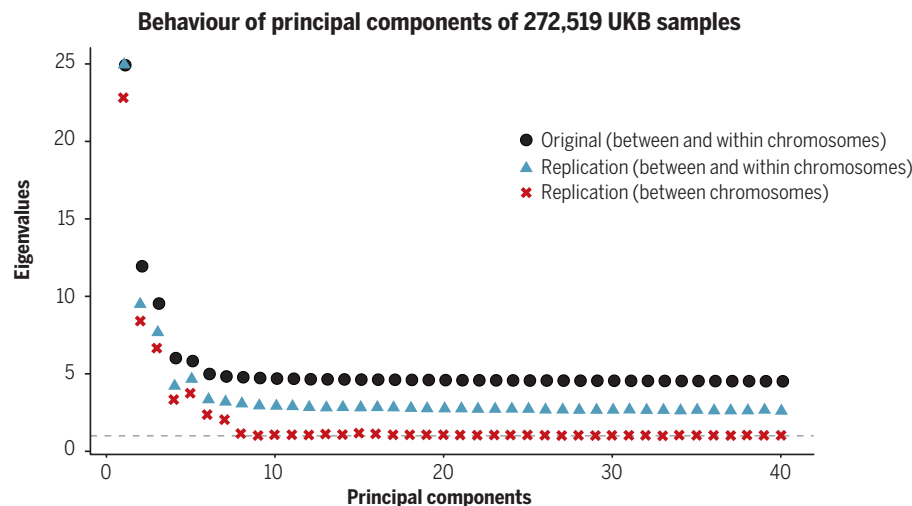


Fig. 3. Behavior of principal components of 272,519 UK Biobank samples. We investigate the degree to which principal components are capturing real population structure by examining whether the genetic variance (eigenvalues) explained by the top 40 principal components inferred from 146,082 SNPs in 272,519 UK Biobank White British (WB) samples replicates in an independent sample of WB. A replication eigenvalue above 1 indicates that the inferred principal component is capturing replicable correlations between SNPs, either local LD (within a chromosome) or population structure (mostly between chromosomes). Original (black circles): eigenvalues of the principal components in the original set of 272,519 WB individuals. Replication (blue triangles): eigenvalue equivalents, i.e., variances of the linear combinations of SNP genotypes using weights inferred from the original set and standardized genotypes in the replication set of 64,969 WB individuals. Replication (between chromosomes only) (red crosses): using the same replication set, but eigenvalue equivalents computed by ignoring the covariances of SNP pairs within the same chromosomes, and counting only the covariances of SNP pairs on different chromosomes, which includes 94.8% of all SNP pairs. The average eigenvalue for the last 32 PCs decreases from 4.37 for the original set to 2.61 for the replication set and further to 1.03 for the between-chromosome set, indicating those PCs are mostly capturing noise and local LD rather than population structure.

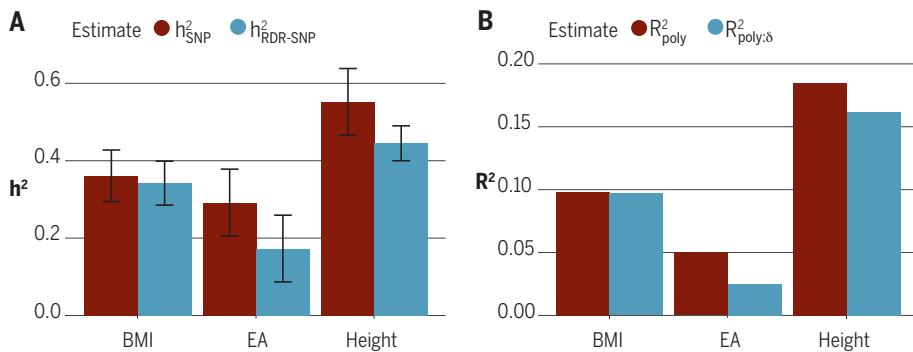


Fig. 4. Shrinkage of polygenic prediction and heritability estimates using within-family designs using Icelandic data. (A) An estimate of SNP heritability using transmitted alleles is given by h_{SNP}^2 ; an estimate of SNP heritability using a within-family method, relatedness disequilibrium regression (RDR) (31), is given by $h_{RDR-SNP}^2$. Statistically significant differences ($P < 0.05$, one-sided z-test) were observed for EA $h_{SNP}^2/h_{RDR-SNP}^2 = 1.72$ ($P = 7.6 \times 10^{-3}$) and height $h_{SNP}^2/h_{RDR-SNP}^2 = 1.24$ ($P = 0.015$). (B) The variance explained by regression of trait onto polygenic score is given by R_{poly}^2 ; the variance explained by a polygenic score when its effect is estimated using a within-family (trio) design is given by $R_{poly;\delta}^2$ (20). We emphasize the relative size of the estimates from within-family methods ($h_{RDR-SNP}^2$ and $R_{poly;\delta}^2$) to between-family methods (h_{SNP}^2 and R_{poly}^2). Between-family methods capture indirect genetic effects from relatives and, potentially, population stratification and assortative mating in addition to the heritability captured by within-family methods. Trait abbreviations: BMI, body mass index; EA, educational attainment (years).

distantly related individuals. The prediction accuracy of PGS is also expected to decrease across ancestry groups because GWAS do not identify causal sites, but sets of possible causal sites in local LD; because local LD patterns depend on population histories, the associations observed in one population will tend to capture causal SNPs less well in others. As expected, recent studies report that the incremental R^2 for a wide range of traits is lower in individuals whose ancestries differ from those of the GWAS set (59, 60).

In addition to allele frequency and LD differences, other factors may contribute to decreased PGS predictive ability: The extent of environmental variance may differ among groups of dissimilar ancestries or selected by distinct enrollment criteria (18), and phenotype measurement may differ across groups. Moreover, effect sizes of variants may differ as a result of gene-gene (GxG) and GxE interactions. Changes in effect sizes may be particularly important for traits to which indirect effects or assortative mating make a large contribution, as such factors could be contingent on cultural and environmental factors. Here, it becomes essential to decompose the nature of the signals identified in GWAS in order to identify which components (e.g., direct versus indirect effects) provide more readily generalized predictions.

Outlook

For many complex traits, GWAS has changed the landscape of genetic investigations and our understanding of genetic architectures. Where once there was not a single reliably replicated association, now there are thousands of variants with robust associations. Notably, GWAS does not require family data, thereby facilitating the collection of large sample sizes. Recently, however, the unique proper-

ties of family data are being brought back to the forefront. For one, some rare variants with strong effects only exist in extended families. Most important, for deeper and more subtle questions, family data such as parent-offspring trios and sib pairs may be necessary to discriminate direct from indirect effects and other confounding factors. Statistically, one natural extension is to extend the study unit from an individual to the nuclear family. In this regard, it is worth noting that as sample sizes increase, close relatives will inevitably be collected as larger fractions of the population are sampled.

A remaining challenge is the issue of ascertainment bias, which occurs when study samples differ systematically from the population. Most sample sets are biased toward individuals of European ancestry (60) as well as toward individuals with higher social economic status and greater health (61), along with other unknown biases. While not necessarily introducing false positives, these ascertainment biases limit the portability of GWAS findings (18, 60). Particularly salient in this regard are GxE interactions, not only over space—that is, across populations at a given time—but over time, given the massive secular trends in environment that have occurred and continue to occur. This consideration applies to health traits, education-related traits, and fertility traits, which affect selection pressures. In this regard, it is important to sample not only different ancestries and current environments but, where possible, to also collect data on multiple generations.

REFERENCES AND NOTES

1. D. Botstein, R. L. White, M. Skolnick, R. W. Davis, *Am. J. Hum. Genet.* **32**, 314–331 (1980).
2. B. Kerem et al., *Science* **245**, 1073–1080 (1989).
3. M. E. MacDonald et al., *Cell* **72**, 971–983 (1993).

4. Y. Miki et al., *Science* **266**, 66–71 (1994).
5. R. Wooster et al., *Nature* **378**, 789–792 (1995).
6. N. Risch, K. Merikangas, *Science* **273**, 1516–1517 (1996).
7. R. J. Klein et al., *Science* **308**, 385–389 (2005).
8. B. M. L. Baselmans et al., *Nat. Genet.* **51**, 445–451 (2019).
9. M. R. Nelson et al., *Nat. Genet.* **47**, 856–860 (2015).
10. Schizophrenia Working Group of the Psychiatric Genomics Consortium, *Nature* **511**, 421–427 (2014).
11. D. J. Schaid, W. Chen, N. B. Larson, *Nat. Rev. Genet.* **19**, 491–504 (2018).
12. T. A. Manolio et al., *Nature* **461**, 747–753 (2009).
13. B. Walsh, M. Lynch, *Evolution and Selection of Quantitative Traits* (Oxford Univ. Press, 2018).
14. E. Marouli et al., *Nature* **542**, 186–190 (2017).
15. L. M. Nolte et al., *Eur. J. Hum. Genet.* **25**, 877–885 (2017).
16. A. V. Khera et al., *Nat. Genet.* **50**, 1219–1224 (2018).
17. A. G. Allegrini et al., *Mol. Psychiatry* **24**, 819–827 (2019).
18. H. Mostafavi, A. Harpak, D. Conley, J. K. Pritchard, M. Przeworski, bioRxiv 629949 [preprint], 7 May 2019.
19. J. N. Hirschhorn, M. J. Daly, *Nat. Rev. Genet.* **6**, 95–108 (2005).
20. A. Kong et al., *Science* **359**, 424–428 (2018).
21. R. J. Hung et al., *Nature* **452**, 633–637 (2008).
22. B. Bulik-Sullivan et al., *Nat. Genet.* **47**, 1236–1241 (2015).
23. A. L. Price et al., *Nat. Genet.* **38**, 904–909 (2006).
24. L. Yengo et al., *Nat. Hum. Behav.* **2**, 948–954 (2018).
25. A. Agrawal, A. M. Chiu, M. Le, E. Halperin, S. Sankararaman, bioRxiv 729202 [preprint], 8 August 2019.
26. J. J. Berg et al., *eLife* **8**, e39725 (2019).
27. J. Yang, N. A. Zaitlen, M. E. Goddard, P. M. Visscher, A. L. Price, *Nat. Genet.* **46**, 100–106 (2014).
28. N. R. Wray, K. E. Kempner, B. J. Hayes, M. E. Goddard, P. M. Visscher, *Genetics* **211**, 1131–1141 (2019).
29. P.-R. Loh et al., *Nat. Genet.* **47**, 284–290 (2015).
30. J. J. Lee et al., *Nat. Genet.* **50**, 1112–1121 (2018).
31. A. I. Young et al., *Nat. Genet.* **50**, 1304–1310 (2018).
32. S. Trejo, B. W. Domingue, bioRxiv 524850 [preprint], 18 January 2019.
33. J. Yang, J. Zeng, M. E. Goddard, N. R. Wray, P. M. Visscher, *Nat. Genet.* **49**, 1304–1310 (2017).
34. J. Yang et al., *Nat. Genet.* **47**, 1114–1120 (2015).
35. P. Wainschtein et al., bioRxiv 588020 [preprint], 25 March 2019.
36. P. M. Visscher et al., *PLOS Genet.* **10**, e1004269 (2014).
37. B. K. Bulik-Sullivan et al., *Nat. Genet.* **47**, 291–295 (2015).
38. W. Astle, D. J. Balding, *Stat. Sci.* **24**, 451–471 (2009).
39. B. Devlin, K. Roeder, *Biometrics* **55**, 997–1004 (2004).
40. H. K. Finucane et al., *Nat. Genet.* **47**, 1228–1235 (2015).
41. P. Turley et al., *Nat. Genet.* **50**, 229–237 (2018).
42. J. Zheng et al., *Curr. Epidemiol. Rep.* **4**, 330–345 (2017).
43. B. F. Voight et al., *Lancet* **380**, 572–580 (2012).
44. L. Y. Millwood et al., *Lancet* **393**, 1831–1842 (2019).
45. B. Brumpton et al., bioRxiv 602516 [preprint], 5 July 2019.
46. G. Hemani, J. Bowden, G. Davey Smith, *Hum. Mol. Genet.* **27**, R195–R208 (2018).
47. G. H. Freeman, *Heredity* **31**, 339–354 (1973).
48. M. Eichelbaum, M. Ingelman-Sundberg, W. E. Evans, *Annu. Rev. Med.* **57**, 119–137 (2006).
49. W. J. Gauderman et al., *Am. J. Epidemiol.* **186**, 762–770 (2017).
50. A. I. Young, F. Wauthier, P. Donnelly, *Nat. Commun.* **7**, 12724 (2016).
51. T. O. Kilpeläinen et al., *PLOS Med.* **8**, e1001116 (2011).
52. S. H. Barcellos, L. S. Carvalho, P. Turley, *Proc. Natl. Acad. Sci. U.S.A.* **115**, E9765–E9772 (2018).
53. J. Tyrrell et al., *Int. J. Epidemiol.* **46**, 559–575 (2017).
54. M. R. Robinson et al., *Nat. Genet.* **49**, 1174–1181 (2017).
55. G. Paré, N. R. Cook, P. M. Ridker, D. I. Chasman, *PLOS Genet.* **6**, e1000981 (2010).
56. A. I. Young, F. L. Wauthier, P. Donnelly, *Nat. Genet.* **50**, 1608–1614 (2018).
57. E. Cox, *Int. Stat. Rev.* **52**, 1–31 (1984).
58. L. Yengo et al., *Hum. Mol. Genet.* **27**, 3641–3649 (2018).
59. A. R. Martin et al., *Am. J. Hum. Genet.* **100**, 635–649 (2017).
60. A. R. Martin et al., *Nat. Genet.* **51**, 584–591 (2019).
61. A. Fry et al., *Am. J. Epidemiol.* **186**, 1026–1034 (2017).

ACKNOWLEDGMENTS

Thanks to D. Conley, A. Harpak, and J. Pritchard for comments on a draft of the manuscript. Supported by the Li Ka Shing Foundation (A.I.Y., S.B., and A.K.) and NIH grant R01 GM121372 (M.P.). This research has been conducted using the UK Biobank Resource under Application Number 11867.

10.1126/science.aax3710



Deconstructing the sources of genotype-phenotype associations in humans

Alexander I. Young, Stefania Benonisdottir, Molly Przeworski, and Augustine Kong

Science **365** (6460), . DOI: 10.1126/science.aax3710

View the article online

<https://www.science.org/doi/10.1126/science.aax3710>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Downloaded from <https://www.science.org> on October 18, 2023

Use of this article is subject to the [Terms of service](#)

Science (ISSN 1095-9203) is published by the American Association for the Advancement of Science. 1200 New York Avenue NW, Washington, DC 20005. The title *Science* is a registered trademark of AAAS.

Copyright © 2019 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works