

Feature Review

Genetic prediction of complex traits with polygenic scores: a statistical review

Ying Ma¹ and Xiang Zhou ^{1,2,*}

Accurate genetic prediction of complex traits can facilitate disease screening, improve early intervention, and aid in the development of personalized medicine. Genetic prediction of complex traits requires the development of statistical methods that can properly model polygenic architecture and construct a polygenic score (PGS). We present a comprehensive review of 46 methods for PGS construction. We connect the majority of these methods through a multiple linear regression framework which can be instrumental for understanding their prediction performance for traits with distinct genetic architectures. We discuss the practical considerations of PGS analysis as well as challenges and future directions of PGS method development. We hope our review serves as a useful reference both for statistical geneticists who develop PGS methods and for data analysts who perform PGS analysis.

PGS analysis for genetic prediction of complex traits

Complex traits are traits that do not perceivably follow simple Mendelian inheritance laws. Examples include binary traits such as type 2 diabetes and hypertension as well as continuous traits such as body mass index and standing height. Complex traits are influenced by multiple genetic factors including genotype, gene expression, epigenomic modifications, and chromatin structure, as well as multiple environmental factors including occupational, lifestyle, and environmental exposures [1,2]. Among these factors, genotypes, in the form of **SNPs** (see [Glossary](#)), represent one of the earliest, most stable, and accurately measurable factors underlying complex traits [3]. In particular, the genotypes of an individual remain the same across somatic cells and tissues over lifetime, with mutations being extremely rare and often neutral [4]. In addition, the genotypes of an individual can be accurately measured in a cost-effective way through various array-based and sequencing-based technologies, and can be further imputed across millions of genomic locations [5–7]. Therefore, genotypes can be used to predict complex traits and reconstruct the genetic predisposition of an individual to a particular disease long before disease onset [8,9]. Such genetic prediction of complex traits can facilitate disease screening and prevention at a population scale, improve symptom diagnosis and intervention at an early stage, and aid in the development of precision medicine with individual-based treatment choices [10–13].

Genetic prediction of complex traits is often carried out by PGS analysis. The PGS for a trait, in its simplest form, is a weighted summation of genotypes across SNPs with the weights being the estimated genetic **effect sizes** [10,14–16]. The PGS is commonly referred to as the polygenic risk score (PRS) or genetic risk score (GRS) when the trait of interest is a binary trait of disease status [13,16,17]. PGS analysis has become increasingly popular ([Figure 1A](#), Key figure) with the abundant availability of genotype and phenotype information collected from various **genome-wide association study (GWAS)** analyses [15,18,19]. In the past decade, GWAS analysis has not only successfully identified many SNPs associated with various complex traits [17,20–22] but also demonstrated that most complex traits have a polygenic [23–25] or

Highlights

Polygenic score (PGS) analysis aggregates association information from genome-wide SNPs to enable genetic prediction of complex traits.

PGS analysis is becoming increasingly popular with the abundant availability of genome-wide association studies and the development of PGS methods.

Different PGS methods model the polygenic architecture underlying traits in different ways and often make distinct modeling assumptions on the effect size distribution. These modeling assumptions can help to understand the performance of PGS methods across traits with distinct genetic architectures.

Recent PGS methods focus on making use of summary statistics as input, specify flexible effect size assumptions, incorporate additional information including SNP functional annotations and pleiotropy association evidence across multiple traits, perform multi-ethnic prediction, and rely on computationally efficient algorithms for scalable inference.

The development of PGS methods is closely connected to the development of methods for SNP heritability estimation, and many common methods are shared between the two areas. Experience and lessons learned from SNP heritability estimation can potentially benefit methodological developments for PGS construction.

For some diseases, PGS analysis has had initial clinical success and can be especially useful for risk stratification. For the majority of complex traits, however, PGS methods have yet to achieve high prediction accuracy in the general population.



omnigenic architecture [26] with an appreciable heritable component. Indeed, many complex traits are influenced by thousands of small-effect SNPs [27,28], which together can explain a substantial proportion of the phenotypic variance, a quantity known as **SNP heritability** [29,30]. Consequently, the prediction of complex traits by using a handful of SNPs that pass the stringent genome-wide significance level is not optimal [20,31]. Instead, genetic prediction of complex traits requires PGS methods that can jointly model genome-wide SNPs.

The development of PGS methods has a long-standing history in both animal breeding programs and human genetics [32]. In animal breeding programs, PGS methods are commonly used for predicting the **breeding values** of animals, which are the expected phenotypic values of an individual's offspring. There, PGS is referred to as the genomic estimated breeding value (GEBV), and PGS-based selective breeding is also referred to as genomic selection. Since the seminal paper of Meuwissen *et al.* [33], genomic selection has achieved remarkable progress in many animal programs and has led to substantial increases in many breeding values such as dairy cattle traits [34]. In human genetics, Wray *et al.* [35] evaluated the feasibility and accuracy of predicting the genetic risk of disease by using dense genome-wide SNP panels. Later, the predicted genetic risk of disease was termed the PRS [24]. For some diseases, PGS methods have had initial clinical success [36,37] and are applied in counseling, prophylactic intervention, and embryonic screening [38–41]. For the majority of common diseases and quantitative traits, the PGS currently has a relatively low overall prediction accuracy across individuals in the general population but can be effective for **risk stratification** that aims to identify individuals with high disease risk [17,18]. In addition, PGS analysis has many other applications beyond phenotype prediction. For example, the PGS for a trait of interest can be treated as a covariate in a **phenome-wide association study (PheWAS)** to identify clinical phenotypes and risk factors that are associated with the genetic predisposition of the trait [42,43]. The PGS can also be viewed as the combined effects of multiple instrumental variables, and is applied in Mendelian randomization analysis to study the causal relationships between complex traits [42,44]. Importantly, PGS accuracy is expected to improve along with increasing GWAS sample size, the availability of new genomic information from omic studies, and the development of advanced PGS methods.

A plethora of PGS methods have already been developed in recent years (Figure 1B–D) [45]. These take advantage of the polygenic architecture underlying complex traits and model it in different ways. We present here a comprehensive review of 46 PGS methods (Table S1 in the supplemental information online), with a primary focus on methods that make use of summary statistics. For completeness and methodological coherence, we have included early individual-level data-based PGS methods, and we therefore introduce PGS methods in non-chronological order. In contrast to previous PGS reviews that focused on the practical interpretation and clinical applications of PGS analysis [11,12,16,17,40], we focus on the methodological aspect of PGS methods and review them from a statistical perspective. In particular, we connect the majority of PGS methods through a multiple linear regression modeling framework and show how different PGS methods can be viewed as making distinct modeling assumptions about the distribution of SNP effect sizes across the genome. We show that such a modeling framework can be instrumental for understanding the behavior and prediction accuracy of different PGS methods for traits with distinct genetic architectures. Based on this framework, we discuss the practical considerations of PGS analysis as well as current challenges and future directions of PGS method development.

A multiple linear regression framework

We begin by introducing a simple multiple linear regression model that relates genotypes to the phenotype of interest. To do so, we define y as a n -vector of the phenotypes measured in n

¹Department of Biostatistics, University of Michigan, Ann Arbor, MI 48109, USA

²Center for Statistical Genetics, University of Michigan, Ann Arbor, MI 48109, USA

*Correspondence: xzhousph@umich.edu (X. Zhou).

individuals by the GWAS. We assume for the present that the phenotype of interest is quantitative, but we discuss the case of binary phenotypes in a later section. We define X as the n by p matrix of genotypes collected across p SNPs on the same set of individuals. Genotypes are often coded as the number of reference alleles for each SNP, and can be represented as continuous values between 0 and 2 after imputation. To simplify discussion, we assume that the phenotype vector y and each column of the genotype matrix X have been centered to have a mean of zero. Centering does not influence the results and allows us to ignore the intercept in the equation. We consider the following multiple linear regression model that relates X to y :

$$y = X\beta + \varepsilon \quad [1]$$

where β is a p -vector of SNP effect sizes, ε is a n -vector of residual errors, and each element follows an independent normal distribution, or $\varepsilon_j \sim N(0, \sigma_\varepsilon^2)$.

Despite its simplicity, the above model is instrumental for understanding almost all existing PGS methods. In particular, most PGS methods can be viewed as making distinct modeling assumptions for the SNP effect size β in the model and rely on different algorithms to obtain the estimates $\hat{\beta}$. The SNP effect size estimates $\hat{\beta}$ subsequently serve as the SNP weights for constructing a PGS for newly observed individuals (Box 1). Because the model includes genome-wide SNPs that are in potential **linkage disequilibrium (LD)** with each other as covariates, the resulting PGS naturally accounts for SNP LD.

Sparse modeling assumptions for SNP effect sizes

One common modeling assumption for the effect size β in the multiple linear regression model is sparsity, and one common sparsity assumption is a point-normal distribution (Box 1). The point-normal distribution assumes that only a small proportion of SNPs have non-zero effects and that their effect sizes follow a normal distribution with a mean of zero and a variance of σ_β^2 . PGS methods that use the point-normal distribution include the Bayesian variable selection regression (BVSr) [46,47], the Bayesian alphabetic method BayesC π [48], LDpred [31], and JAMPred [49]. The first two methods use individual-level data from GWAS analysis, whereas the last two use GWAS summary statistics. These sparse PGS methods also have subtle differences in their assumptions regarding the hyper-parameters as well as their Markov chain Monte Carlo (MCMC) fitting algorithms.

The point-normal distribution assumes that the effect sizes of the non-zero effect SNPs follow a normal distribution. The normality assumption for SNP effect sizes is highly effective in many genetic applications including SNP heritability estimation [20,30], and is commonly referred to as the global shrinkage approach [50]. However, the normality assumption has a potential drawback: the normal distribution has a thin tail, which corresponds to a relatively low prior probability of observing large effect sizes. Consequently, normality assumption can lead to over-shrinkage of large effect estimates that are crucial for accurate prediction. Because of the drawback in the normality assumption, several PGS methods have been developed to introduce heavy tailed distributions of the non-zero effects to ensure adaptive shrinkage, also known as local shrinkage [50]. These methods often assume an SNP-specific non-zero effect size variance σ_j^2 for the SNP j and place another prior distribution on σ_j^2 . The prior on σ_j^2 can be either continuous or discrete, effectively leading to a scale-mixture of normal distribution on the non-zero effect sizes. For example, BayesB [33,48] places an inverse gamma (IG) distribution on σ_j^2 , leading to a point- t distribution for β . BayesD [48] and BayesD π [48] place a mixture of IG distributions on σ_j^2 , leading to a point- t mixture distributions for β . BayesR [51] places a discrete distribution on σ_j^2 , leading to a mixture of

Glossary

Best linear unbiased prediction

(BLUP): this is used in linear mixed models for estimating and predicting the random effects. BLUP is a linear function of the outcome, is an unbiased predictor of the random effects, and is best in the sense that the variance of the prediction error, in the form of the mean squared difference between the estimated values and truth, is not greater than that obtained from any other linear unbiased predictors. The BLUPs of random effects are similar to the best linear unbiased estimates (BLUEs) of fixed effects.

Breeding values: the expected phenotypic values of the offspring of an individual.

Clumping: the procedure of selecting a subset of SNPs that are approximately independent of each other.

Effect size: the coefficient of a SNP genotype for an outcome phenotype of interest. It is closely related to the proportion of phenotypic variance that is explained by the genotype.

Genome-wide association study

(GWAS): an experimental design that aims to identify SNPs or other genetic variations associated with traits of interest based on samples collected from populations.

Linkage disequilibrium (LD): the phenomenon that two alleles at different loci occur together in the same gamete more often than would be expected by chance alone. The coefficient of LD is defined as the difference between the frequency of gametes carrying the pair of two alleles at two loci and the product of the frequencies of those two alleles. For PGS studies, LD is often calculated as the correlation between SNP genotypes using potentially unphased genotype data.

Phenome-wide association study

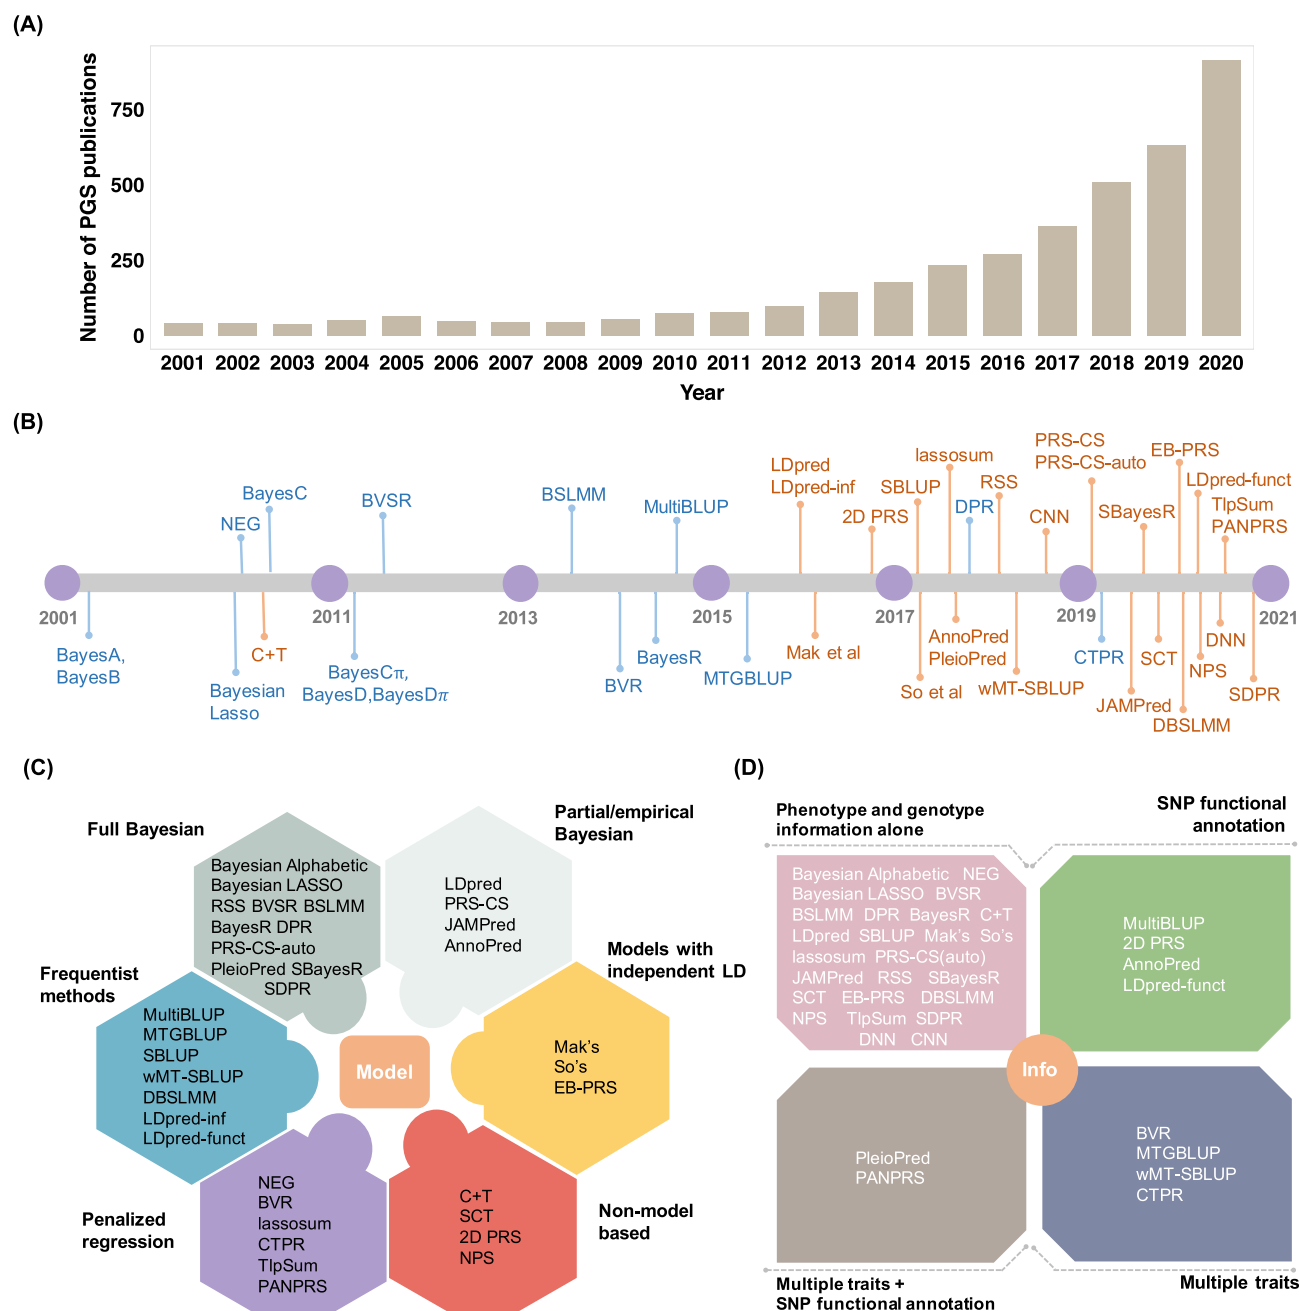
(PheWAS): a study that focuses on identifying phenotypes associated with a covariate of interest, which is often a genetic variant or the PGS of another phenotype.

Restricted maximum likelihood

(REML): a particular form of maximum likelihood estimation procedure for linear mixed models that produces unbiased estimates for variance and covariance parameters. It is based on a likelihood function defined on a restricted subset of parameters after integrating out the nuisance parameters.

Key figure

An overview of polygenic score (PGS) methods



Trends in Genetics

(See figure legend at the bottom of the next page.)

Box 1. Predicting new observations through PGS construction

We can predict phenotypes for newly observed individuals using the estimated SNP effect size $\hat{\beta}$ from the model in Equation 1 (main text). Specifically, once we have obtained the p -vector of genotypes, x_i , for a new individual i we can simply plug in the SNP effect estimates to obtain the predicted phenotype value (i.e., the PGS), as $\hat{y}_i = x_i \hat{\beta}$.

Common modeling assumptions for SNP effect sizes

Because $p \gg n$, we will need to make additional modeling assumptions for the effect size β to make the model in Equation 1 (main text) identifiable. Both sparse and polygenic modeling assumptions have been proposed for β . A common sparse modeling assumption for β is the point-normal distribution, which assumes that the effect size of SNP j comes from a mixture of a normal distribution and a point mass at zero, expressed as:

$$\beta_j \sim \pi N(0, \sigma_\beta^2) + (1 - \pi) \delta_0 \quad [I]$$

where, with proportion π , the SNP effect size follows a normal distribution with mean zero and variance σ_β^2 , and with a proportion of $1 - \pi$ the effect size is exactly zero – hence the point mass at zero, δ_0 . In the point-normal distribution, π is usually assumed to be small, representing the prior belief that a small proportion of SNPs have non-zero effects.

A common polygenic modeling assumption for β is the normal assumption, which assumes that all SNPs have non-zero effects and that each effect size follows a normal distribution:

$$\beta_j \sim N(0, \sigma_\beta^2) \quad [II]$$

with mean zero and variance σ_β^2 . The model in Equation 1 (main text), when paired with the normality assumption for β in Equation II, has a wide variety of applications and many names. For example, it is referred to as the linear mixed model (LMM) per the resulting random effects term of the combined genetic effects; as the infinitesimal model per its polygenic assumption on the effect sizes $X\beta$; as the ridge regression in statistics literature; as the L2 regularization when viewed as a penalized regression; as the **best linear unbiased prediction (BLUP)** when the focus is on the predicted values; or as the REML when the REML algorithm is used for inference. All these names are used interchangeably in the PGS literature, and we simply refer to the model as the LMM in the present review.

Risk stratification: the procedure of systemically categorizing individuals into subgroups based on their predicted risks, with a special emphasis on identifying individuals with a particularly high disease risk for optimized medical decision making. Risk stratification is conceptually different from risk prediction which aims to predict disease risk for all individuals in a population.

SNP: the most common type of genetic variation at a single position in the DNA sequence. A SNP occurs when a single nucleotide in the genome differs between individuals or between paired chromosomes in an individual.

SNP heritability: the proportion of phenotypic variance in the outcome trait that is explained by measured SNP genotypes in a GWAS. Usually only additive genetic factors are considered.

three normal distributions together with a point mass at zero for β . BayesR is further extended by SBayesR [52] to take summary statistics as the input. All these methods rely on MCMC for model inference.

The above PGS methods make explicit sparse modeling assumptions to induce sparsity for effect size estimates. Several PGS methods that were initially described as an algorithm can also be viewed as making implicit sparse modeling assumptions. For example, the most commonly used PGS method, C+T [24,35,53], relies on LD **clumping** and P value thresholding to select a subset of approximately independent SNPs with strong association signals for PGS construction. The C+T strategy ensures that a sparse set of SNPs is used for constructing the PGS, and thus corresponds to making a sparse assumption for SNP effect sizes. Similarly, SCT [54] extends C+T by examining an extended set of hyper-parameters for SNP selection. These hyper-parameters include P value threshold, LD window size, LD correlation threshold, and

Figure 1. (A) The number of PGS publications increased substantially from 2001 to 2020, highlighting the popularity of PGS analysis. The number of publications was obtained by searching using the key terms 'polygenic + score + or + polygenic + risk + score' on PubMed. (B) Timeline of commonly used PGS methods developed over the past two decades. These PGS methods either use individual-level genotype and phenotype data as the input (blue) or use summary statistics as the input (orange). (C) PGS methods can be categorized into six categories based on their model and fitting strategy. Specifically, some PGS methods are model-based and are described as a formal model with a corresponding fitting algorithm (colors other than red), whereas others are algorithm-based and are described as an algorithm or a fitting procedure without an explicit model (red). The model-based PGS methods can be further categorized based on the underlying inference algorithm: some are fully Bayesian and use Markov chain Monte Carlo (MCMC) for model fitting (grey); some are partial/empirical Bayesian that optimize particular hyper-parameters through grid search while obtaining other parameter estimates through MCMC (light grey); some are approximate approaches that assume independence across SNPs and use optimization for effect size estimation (yellow); some are frequentist in nature and can obtain an analytic solution without optimization (blue); and some are based on penalized regression and use iterative algorithms for parameter estimation (purple). (D) PGS methods can also be categorized in terms of the information used for PGS construction. Most PGS methods use only genotype and phenotype information from the GWAS on the trait of interest (pink). Some recent PGS methods can use additional SNP annotation information obtained from external data sources (green) and/or other phenotype information in addition to the phenotype of interest (taupe and navy blue).

imputation score. PGS values in SCT are constructed for different combination of the hyper-parameters and are further selected through penalized regression in the validation data.

Polygenic modeling assumptions for SNP effect sizes

An alternative to the sparse modeling assumption for the effect sizes is the polygenic modeling assumption, and the most common polygenic modeling assumption is normality (Box 1). The model in Equation 1, when paired with the normality assumption for β in Equation II in Box 1, has a wide variety of applications and has many names (Box 1). We refer here to the model simply as the linear mixed model (LMM), which has been implemented in many software. For example, GEMMA [47] implements LMM for prediction using individual-level data. LDpred-inf [31], SBLUP [55], and DBSLMM [56] all implement the same model using summary statistics as the input.

Similarly to sparse modeling, multiple PGS methods have been proposed to extend the normality assumption in the polygenic models to enable more accurate prediction. For example, BayesA [25,33,57] places an IG distribution on the SNP specific variance σ_j^2 to induce a t distribution for the effect size β . NEG [58] places an exponential-gamma distribution on σ_j^2 to induce a normal-exponential-gamma distribution for β . PRS-CS [59] decomposes σ_j^2 as a product of two parameters: a global shrinkage parameter either placed with a half Cauchy prior or optimized through a grid search, and a local shrinkage parameter with a gamma-gamma prior. Both BayesA and NEG use individual-level data as the input whereas PRS-CS uses summary statistics. As another popular example, the Bayesian version of Lasso [60] effectively places an exponential distribution on σ_j^2 to induce a Laplacian/double exponential prior for β . The Bayesian Lasso is fitted through either MCMC [61,62] or the EM algorithm [63] to obtain the posterior mean of β . By contrast, the frequentist Lasso is expressed in the form of an L1 penalized regression and is often fitted through a gradient descent algorithm to effectively obtain the posterior mode for β . Although the posterior mean of β is not sparse, the posterior mode is. For PGS construction, the lassosum [64] fits the frequentist Lasso using summary statistics as the input. TipSum [65] extends lassosum by selectively penalizing small effect SNPs via the truncated Lasso penalty.

In addition to placing a continuous prior on σ_j^2 , several PGS methods also place a discrete prior on σ_j^2 to effectively induce a mixture of normal distributions for β . For example, the Bayesian sparse linear mixed model (BSLMM) [66] assumes that each effect size comes from a mixture of two normal distributions. By segregating SNPs into two categories, BSLMM can place different shrinkages on the SNP effect sizes in the two categories separately, leading to proper shrinkage of small effects without over-shrinkage of large effects. BSLMM is implemented in GEMMA [47], which takes individual-level data as the input and relies on MCMC for inference. BSLMM is also implemented in DBSLMM [56], which takes summary statistics and relies on an efficient deterministic algorithm for scalable inference. As another example, BayesC [57,67] places a mixture of IG distributions on σ_j^2 , thus inducing a mixture of t distributions for β .

The two types of polygenic extensions to normality based on continuous and discrete priors on σ_j^2 have different modeling benefits. Specifically, a continuous prior on σ_j^2 often leads to an effect size distribution that is relatively easy to perform inference on, whereas a discrete distribution on σ_j^2 often allows more adaptive shrinkage of the effect sizes and robust prediction performance across traits. A common feature of both extensions is that they are parametric in nature, relying on the use of a limited number of parameters to characterize the effect size distribution, which can be restrictive. To enable more flexible effect size modeling, the latent Dirichlet process regression (DPR) [68] uses a Bayesian non-parametric model and places a distribution on σ_j^2 , where the distribution is inferred based on the data at hand. The non-parametric distribution on σ_j^2 in DPR

leads to a normal mixture with infinitely many components of the effect sizes, making DPR robust and adaptive to a wide range of phenotypes with different genetic architectures. DPR is implemented in the DPR package that uses individual-level data as the input and relies on either MCMC or variational Bayes for inference. DPR is also implemented in SDPR [69], which takes summary statistics as the input.

The above PGS methods make explicit polygenic modeling assumptions. A few PGS methods that were originally described as a fitting algorithm can also be viewed as making implicit polygenic modeling assumptions. For example, Mak *et al.* [70] constructed the PGS by weighting SNP marginal effect size estimates using local true discovery rates that are estimated through either maximum likelihood or kernel density estimation. Because the local true discovery rate ranges between 0 and 1, the Mak *et al.* method implicitly assumes that all SNPs are included for PGS construction. So and Sham [71] extend the approach of Mak *et al.* by applying a Tweedie formula [72] to further correct for the SNP effect size estimates before weighting.

Modeling assumptions and other factors that influence performance

Given that the majority of PGS methods make distinct modeling assumptions about the effect sizes, one naturally wonders which PGS method to choose for a given trait. Intuitively, if the prior effect size distribution closely matches the true effect size distribution underlying the trait, then the inferred effect size estimates would approximate well the underlying polygenic architecture, thus leading to accurate prediction performance. Indeed, it has been shown that polygenic PGS methods often perform well for polygenic traits [24,66,73,74], whereas sparse PGS methods often perform well for traits in which a small proportion of SNPs have moderate or large effect sizes [26,59,66]. Because the genetic architecture underlying a trait is often unknown and varies across traits [75], it is often beneficial to use a PGS method with a flexible modeling assumption that can adaptively approximate the true effect size distribution across a range of traits. For example, methods that rely on a mixture of normal distributions (e.g., BSLMM, BayesR, DPR) for adaptive modeling of effect sizes often outperform standard LMM that assumes a single normal distribution.

Certainly, how well the effect size assumption matches the underlying truth is only one modeling factor, albeit a major one, that determines prediction performance. Other modeling factors, such as the choice of inference algorithms and the inference strategies regarding the hyper-parameters, can also substantially impact on prediction performance. Specifically, given the same model and sufficient computational resources, exact inference algorithms often outperform approximate ones. For example, the MCMC algorithm for DPR outperforms the variational Bayesian approximation of DPR across traits. However, with limited computational resources, an approximate inference algorithm may become the only viable option. For example, DBSLMM relies on an approximate deterministic algorithm to perform inference on BSLMM, and is much more scalable than the original MCMC algorithm for fitting BSLMM. In addition, it is generally beneficial to infer various hyper-parameters in the model rather than fixing them to certain pre-assigned values. For example, whereas both BVSR and BayesC π fit a similar sparse model, BVSR often outperforms BayesC π by inferring the hyper-parameters instead of fixing them to a prior set of values. The ability to use a large number of parameters and explore a large parameter space can also help with prediction performance. For example, SCT outperforms C+T by performing SNP selection with additional criteria and exploring a larger hyper-parameter space. Fitting algorithms that use individual-level data as the input usually have higher prediction performance than algorithms that take summary statistics because the latter must approximate the LD matrix (more below). However, owing to LD matrix approximation, algorithms using summary statistics are often much more computationally scalable than those using individual-level data. In addition to the above modeling factors, PGS performance also depends on the quality of the input

data [45], the GWAS sample size, and trait SNP heritability [76,77], which represents the potential performance upper limit for PGS [78].

Finally, multiple factors also influence the computational cost of different PGS methods. For example, PGS methods based on global shrinkage of LMM are often faster than PGS methods with local shrinkage or sparsity inducing priors because the former can be fitted based on an analytic solution. For the same model, approximate inference algorithms are computationally faster and use less memory than exact inference algorithms. At the extreme, algorithm-based PGS methods such as CT and SCT are generally more computationally scalable than model-based PGS methods that specify explicit effect size priors and perform formal inference. In addition, PGS methods that rely on summary statistics as the input make explicit approximations for the LD matrix, which can alleviate much of the computational burden associated with modeling of SNP correlation. Software implementation, the use of multithreading or parallel computing environment, and the choice of computational language can also influence the computational cost of PGS methods.

Adaptation of PGS methods towards using summary statistics

Whereas early PGS methods use individual-level genotype and phenotype as inputs, a growing number of PGS methods can make use of summary statistics or are specifically designed to do so. Fitting with summary statistics requires LD approximation, which can lead to reduced accuracy as compared to fitting with individual-level data using the same model [79]. However, fitting with summary statistics can take advantage of the easily accessible summary statistics from various GWAS analyses without privacy concerns and logistic hurdles [18,19], and can lead to substantial computational gains through LD approximation. Therefore, summary statistics-based PGS methods facilitate PGS applications towards large-scale data, which is crucial for ensuring accurate prediction performance.

There are two general strategies for fitting PGS models using summary statistics, with subtle methodological differences between them. The first strategy is to formulate the model with individual-level data and derive the inference algorithm using summary statistics, whereas the second strategy is to model summary statistics directly (Box 2). Both strategies require two

Box 2. Modeling summary statistics

There are two general strategies for fitting PGS models using summary statistics. The first is to formulate the model with individual-level data and derive the inference algorithm using summary statistics. Specifically, the likelihood for the model in Equation 1 (main text) can be expressed as a function of two terms: $X^T y$ and $X^T X$. Subsequently, instead of using individual-level data X and y as inputs for modeling fitting, one only needs to obtain these two forms of summary statistics. $X^T y$ can be obtained through the p -vector of marginal z -scores which is equivalent to the marginal effect size estimates $\hat{\beta}$ when both the phenotype and the genotypes for each SNP are standardized to have mean zero and unit standard deviation. In that case, the z -scores are in the form of $\mathbf{z} = \frac{X^T y}{N}$ when SNP effect sizes are small, where N is the GWAS sample size. $X^T X$ for the standardized genotype matrix can be obtained through a p by p SNP correlation matrix $\mathbf{D} = \frac{X^T X}{N}$, which is also referred to as the LD matrix. With \mathbf{z} and \mathbf{D} as inputs, likelihood-based inference can be carried out as if individual-level data are available.

The second strategy for fitting PGS models with summary statistics is to model summary statistics directly. For example, regression with summary statistics (RSS) models the marginal effect size estimate $\hat{\beta}$ as a function of the underlying effect size β in the form:

$$\hat{\beta} | \beta \sim N(\mathbf{D}\beta, \sigma_e^2 \mathbf{D}) \quad [1]$$

where $\mathbf{D} = \frac{X^T X}{N}$, which is also referred to as the LD matrix, and σ_e^2 is the same error variance as in Equation 1 (main text). The conditional likelihood of β given the hyper-parameters (e.g., σ_e^2) based on Equation 1 is the same as the conditional likelihood of β based on Equation 1 (main text). Therefore, if the hyper-parameters are known, both strategies for fitting PGS models with summary statistics lead to the same likelihood for β . The likelihood for the hyper-parameters based on Equation 1, however, is different from that based on Equation 1 (main text). Note that a more complex form of RSS is available [80] when the SNP genotypes are not standardized.

forms of summary statistics as the input: the p -vector of marginal z-scores and the p by p SNP correlation matrix D , which is also known as the LD matrix. The input z can be easily obtained through simple linear regression in the original GWAS whereas the input D is often estimated from a reference panel of individuals of the same ethnicity (e.g., from the 1000 Genomes Project). Because of the relatively small sample size in the reference panel, the estimated D requires further regularization and approximation to ensure numerical stability for PGS inference. Some PGS methods approximate D with a block-diagonal matrix computed either based on LD [31,56,59,64] or through index-sorting [69], sometimes further adjusted for cross-block correlations caused by long-range LD [49]; some approximate D with a banded matrix based on a sliding window for LD computation [31,55]; some shrink D towards a diagonal matrix with $D = \Lambda D + (1_{n \times n} - \Lambda)/l$, where each element λ_{ij} in Λ is a function of recombination rate between SNPs i and j [52,80]; and some approximate D with a sparse matrix by setting small matrix elements to zero [81]. Regardless of the form of the estimation, a match between the estimated D from the reference panel and the true D in the study sample is crucial to ensure accurate prediction performance [56,82,83].

Incorporating additional information to improve prediction

Several recent PGS methods have been developed to incorporate additional and external information beyond what is available in the GWAS data. Such external information can be either in the form of SNP functional annotations or in the form of other phenotypes in addition to the phenotype of interest. Incorporating external information often improves the accuracy of PGS.

Incorporating SNP functional annotations

Functional annotations for a given SNP are continuous or binary annotations that characterize the functional importance of the genetic variant [84–86]. SNP functional annotations are obtained through functional genomic studies [87–91] and can serve as crucial predictors for SNP effects. For example, SNPs with particular functional annotations are more likely to be causal [92], tend to have larger effect sizes, and explain more heritability than SNPs with other annotations [93,94]. Several PGS methods have been developed to incorporate SNP functional annotations to improve prediction. For example, 2D PRS [95] categorizes SNPs into two disjoint sets: one containing high-priority SNPs that are likely to be associated with the trait of interest, and the other containing low-priority SNPs that are less likely to be associated with the trait. The two sets of SNPs are determined based on a separate GWAS and are then subject to the C+T procedure separately with a less stringent P value threshold for the high-priority SNPs. MultiBLUP [96] divides SNPs into separate groups based on their genomic location and induces different effect size shrinkage in different groups. AnnoPred [97] incorporates SNP functional annotations directly into the prior distribution of effect sizes based on BVSR: it either models the probability that the SNP j has a non-zero effect as a function of its annotations, or models its non-zero effect variance as a function of its annotations. LDpred-funct [98] builds upon LMM and models σ_j^2 as a function of its annotations.

Modeling pleiotropy across multiple traits

Another type of external information that can facilitate trait prediction is pleiotropic association information. Pleiotropic association characterizes SNP effect similarity across multiple correlated traits and can be used to improve SNP effect size estimation for the trait of interest [99–101]. PGS methods that take advantage of pleiotropy are often based on the multivariate linear mixed model (mvLMM) [102,103], also known as the MT-BLUP in prediction settings. The mvLMM is an extension of LMM and assumes that the effects of SNP j across phenotypes follow a multivariate normal distribution, and use a covariance matrix to capture the genetic covariance across traits. By jointly modeling SNP effect size similarity across traits, mvLMM can borrow

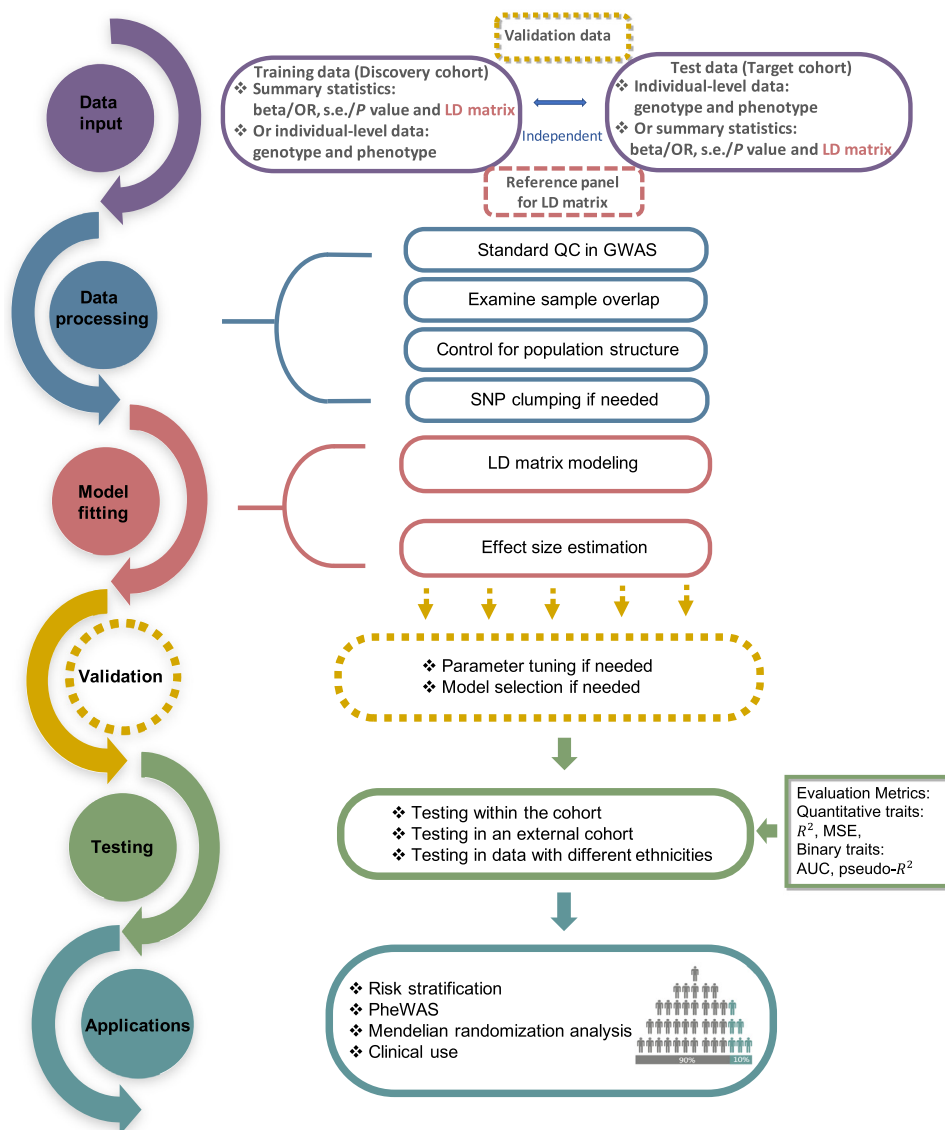
information of effect size estimates from other traits to improve the estimates for the trait of interest. Li *et al.* [104] implement a bivariate version of mvLMM that models two phenotypes jointly. MTGBLUP [100] implemented a general form of mvLMM with individual-level data as input. wMT-SBLUP [99] implements mvLMM with summary statistics as the input. In addition to mvLMM, CTPR [105] imposes a sparse effect size assumption on each trait and uses an L2 penalty to model effect size similarity across traits. Other methods also incorporate SNP functional annotations into pleiotropic modeling. For example, PleioPred [101] partitions SNPs into multiple annotation categories while jointly modeling two correlated traits together. PANPRS [106] specifies an annotation specific L1 penalty for SNPs in each annotation category to incorporate annotation into prediction, and uses a group Lasso-type penalty to encourage SNP effect size similarity across traits.

Moving beyond multiple linear regression

Although the multiple linear regression framework in Equation 1 includes the majority of PGS methods, there are several notable exceptions. For example, the non-parametric shrinkage (NPS) method [81] performs a linear transformation of the SNP genotypes before placing a non-parametric effect size distribution on the transformed genotypes. Subsequently, the resulting prior distribution of the original genotypes from NPS is not straightforward to characterize and does not directly correspond to a known distribution. As another example, deep-learning methods [107,108] rely on deep convolutional neural networks connected through the leaky rectified linear unit (ReLU) activation functions for modeling non-linear effects, which can be particularly effective for predicting traits with appreciable genetic heterogeneity [107]. However, the performance of deep-learning methods is heavily dependent on the availability of large-scale training data, the choice of network architecture, and tuning of hyper-parameters; the latter two require expertise and extensive trial and error owing to the lack of a standard theory to guide architecture selection and model training. For case-control studies, the binary case-control labels are often treated as continuous traits and directly modeled through the multiple linear regression framework [66,68]. Such modeling could be justified by recognizing the linear model as a first-order Taylor approximation to a generalized linear model [66,68]. However, several recent PGS methods use either a logistic regression [106], its approximation [109], or a liability threshold model [66] to directly model ascertainment and the binary nature of the case-control outcome. Finally, recent studies have started to explore the development of PGS methods to predict the absolute risk that an individual develops a disease over a given period of time by using the Cox proportional hazard model [110]. Validating such absolute risk models in prospective studies will be of particular clinical importance [40].

Evaluating PGS methods: cross-validation and cross-ethnicity performance

Fitting and evaluating PGS methods rely on a multistage procedure commonly referred to as cross-validation (Figure 2). Cross-validation requires two or three datasets: training data, optional validation data, and test data. PGS methods are fitted to the training data; if needed, their hyper-parameters are determined from the validation data (Table S1 in the supplemental information online); and eventually their performance is evaluated using the test data. The relative sizes of the training versus test data represent a bias-variance trade-off in estimating the prediction error. In particular, a small training dataset and a large test dataset would likely lead to an overestimate of the prediction error. A large training dataset and a small test dataset, by contrast, would result in less bias but higher variance in estimating the prediction error. In addition, methods that perform automatic inference for all parameters using the training data alone can potentially combine the validation data with the training data to benefit from the larger sample size. By contrast, methods that tune hyper-parameters in a separate validation dataset are often computationally easier to fit, and require estimation of the SNP effect sizes conditional on



Trends in Genetics

Figure 2. A general pipeline for polygenic score (PGS) construction and applications. PGS methods require either two or three datasets as the input: training data, test data, and (if necessary) validation data. These datasets need to undergo multiple steps of stringent quality control that include SNP filtering, overlap sample removal, adjustment of population stratification, etc. The training data are then used to fit the desired PGS model for estimating the SNP effect sizes. For some PGS methods, validation data are necessary to tune parameters in the model or perform model selection. The estimated SNP effect sizes are then used to construct PGS values from the test data, where the predictive performance of the PGS method is tested based on standard metrics. The constructed PGS values are used for different applications, including risk stratification, phenome-wide association study (PheWAS), and Mendelian randomization. A dotted line box represents a step that is not necessary for all PGS methods. Abbreviations: AUC, area under the curve; LD, linkage disequilibrium; MSE, mean squared error; OR, odds ratio; PheWAS, phenome-wide association study; QC, quality control; s.e., standard error.

the hyper-parameters in the training data instead of jointly estimating both, although their performance may also be influenced by the size of the validation data. In cross-validation, the evaluation metrics in the test data include R^2 and mean squared error (MSE) for quantitative traits, and area

under the curve (AUC) and pseudo- R^2 for binary traits. Among these metrics, AUC and R^2 are easier to interpret because both range between zero and one, but neither accounts for the predicted trait mean as MSE does, and thus may not be suited for settings where predicting the absolute trait value is of interest. Importantly, tuning of hyper-parameters in the validation data may only require summary statistics, as does computing R^2 [56,65,99] or AUC [111] in the test data. Using summary statistics for hyper-parameter tuning and R^2 computation facilitates the application of PGS methods to a wide variety of datasets [56]. Finally, we note that one unfortunate mistake that practitioners commonly make in the cross-validation procedure is to use the test data instead of a separate validation dataset to tune the hyper-parameters. Using the same test data to both tune hyper-parameters and evaluate PGS performance would lead to model over-fitting, resulting in underestimation of the prediction error.

Most cross-validations have thus far been performed on a single GWAS of samples of European ancestry [112]. Several recent studies have explored PGS evaluation either through cross-study validation where the training and test data are from two separate GWASs, or through cross-ethnic group validation where the training and test data are from two GWASs with samples of different ethnicity [56]. Cross-study and cross-ethnicity PGS applications are challenging because of the potential mismatch in allele frequency and LD pattern between the training and test data [81]. Indeed, models trained with European individuals are often two- to threefold less accurate when applied to Asian or African populations as compared to European populations [56,113,114]. Consequently, special PGS methods have been developed to enhance cross-ethnicity prediction. For example, a weighted multi-ethnic PGS is proposed to combine PGS trained in Europeans and non-Europeans to improve prediction in both populations [115]. PolyPred and PolyPred+ [116] rely on functionally informed fine-mapping in different populations to improve causal effect estimation and subsequent cross-ethnicity prediction accuracy. PRS-CSx [117], an extension of PRS-CS, directly assumes shared causal effects and borrows information across populations for accurate effect size estimation. With methodological advances and increased availability of GWASs in under-represented populations [112,114], PGS accuracy in diverse populations will be further improved.

Concluding remarks

We summarize the discussed PGS methods in a reference guide to facilitate practical applications (Figure 3). The performance of different PGS methods has been evaluated in multiple human traits in both PGS method studies (Figure 4 and Figure S1 in the supplemental information online) and method comparison studies [83,118,119]. These studies have shown that C+T is the most commonly compared method owing to its simplicity and computational efficiency, whereas PRS-CS and BSLMM tend to have higher performance than the others whenever they are compared, presumably because of their flexible modeling assumptions. However, these studies have also shown that different PGS methods have distinct performance across traits and that the same method may have different performance on the same trait in different studies owing to varying cross-validation designs. Therefore, comprehensive comparisons will be necessary to systematically evaluate the performance of various PGS methods in the future.

We note that the development of PGS methods is closely connected to the development of methods for SNP heritability estimation, and many common methods are shared between the two areas [30]. For example, the sparse PGS methods BVSr [46] and BayesR [51], as well as the polygenic PGS methods LMM [20], BSLMM [66], and DPR [68], are all commonly used for SNP heritability estimation. Among them, LMM is perhaps the most widely applied [20], and has multiple software implementations [47,120] and multiple available fitting algorithms including **restricted maximum likelihood (REML)** and the method of moments for SNP heritability

Outstanding questions

What is the best approach to borrow information across multiple ethnic groups to improve the portability of the PGS across different ethnicities while maintaining its accuracy in specific ethnic groups?

What is the best way to approximate the LD matrix such that we can maintain the accuracy of individual-level based PGS methods while keeping the computation benefits from summary statistics-based PGS methods?

Would modeling the SNP effect size dependence on the minor allele frequency and LD help to improve PGS accuracy?

Can we incorporate other integrative approaches recently developed in various omic studies into PGS modeling to improve prediction performance?

Would selecting informative functional annotations and/or selecting correlated traits from a large group of candidates help to further improve PGS performance for the trait of interest?

Can we measure prediction uncertainty through the predictive posterior distribution in a computationally efficient fashion, and can we quantify the calibration of such prediction uncertainty through posterior predictive checks?

How do we extend the current PGS methods to predict the absolute risk that a person develops a disease over a given period of time?

How do we appropriately communicate PGS results, especially their relatively low accuracy in the general population, to patients and consumers who obtained their PGS through clinical and laboratory tests?

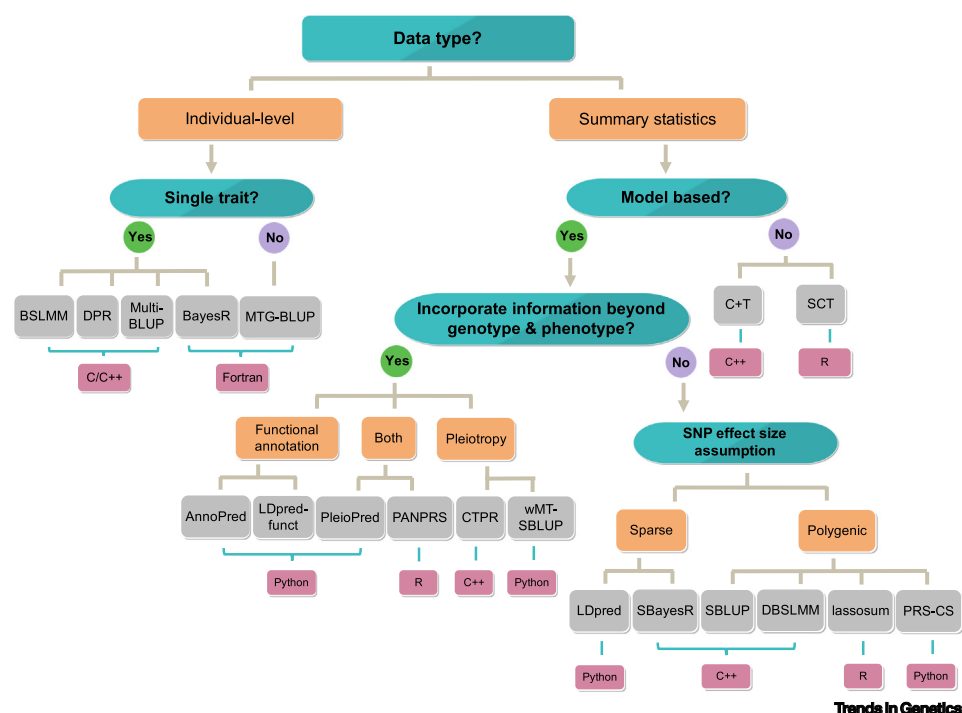


Figure 3. A decision tree for which methods to use for polygenic score (PGS) analysis. The decision tree begins with input data type, followed by the choice of analyzing single versus multiple traits, using model-based methods versus algorithm-based methods, whether to incorporate information beyond genotype and phenotype, as well as the detailed SNP effect size assumptions (blue brackets). The choices include Yes/No answers (Yes in green circles and No in purple circles) or other qualitative options (orange brackets). Different choices lead to different PGS methods (grey brackets) which are implemented with different computing languages (pink brackets).

estimation [121]. In addition to analyzing a single quantitative trait, LMM has also been extended to SNP heritability estimation for binary [66, 122, 123] and count [124–126] traits, as well as for genetic and environmental covariance estimation across multiple phenotypes [103, 127]. With the same model, PGS methods focus on estimating SNP effect sizes, whereas heritability estimation methods focus on estimating a variance component hyper-parameter that represents SNP heritability. The estimated SNP heritability depicts a potential upper limit of PGS performance and serves as an initial input for many PGS methods [56]. As with PGS methods, the accuracy of SNP heritability estimation is highly dependent on how well the prior effect size distribution matches the truth [66]. Indeed, a similar trend in SNP heritability estimation is to develop methods with flexible SNP effect size distributional assumptions, often by incorporating SNP annotations or by modeling the SNP effect size dependence from the minor allele frequency (MAF) and LD score [121, 128]. For example, LDAK assumes that the variance of SNP effect size is a function of MAF and LD, whereas genome-based restricted maximum likelihood (GREML)-MS [129] and stratified LDSC [128] induce such dependence by stratifying SNPs into different MAF and LD bins and assuming different per-SNP heritability values in different strata. Finally, several SNP heritability estimation methods have been developed to take GWAS summary statistics as the input. These summary statistics-based methods include LDSC [130] and MQS [121] algorithms for LMM, and the SumHer algorithm [131] for LDAK, all of which rely on the method of moments to achieve scalable computation. A recent review on SNP heritability estimation from a statistical perspective is available [30]. Taking advantage of the methods developed for SNP heritability estimation, and incorporating the lessons and experiences gained in that research area, can potentially benefit the development of PGS methods.




Figure 4. Predictive performance of common polygenic score (PGS) methods as revealed in the PGS methodological publications. (A) The bar plot shows the top five PGS methods that have been most compared in real data applications in the 26 PGS methodological publications listed in Figure S1 in the supplemental information online. The y axis denotes the number of times a specific PGS method is compared in a different PGS methodological publication. Note that PGS methods developed earlier tend to be compared more often than methods developed later. (B) The bar plot shows the percentage of times a PGS method is ranked among the top two methods in terms of prediction performance in human traits in the PGS methodological publications. The percentage is calculated both across publications and across traits examined in all PGS methodological publications listed in Figure S1 in the supplemental information online. In both (A) and (B) we only considered PGS methods that have been compared at least once in a PGS methodological publication from a different research group.

Although existing PGS methods have shown promising performance across many complex traits, many future improvements are warranted (see [Outstanding questions](#)). For example, annotation-facilitated PGS methods have so far focused on a limited number and types of annotations. Evaluating a large variety of annotations and exploring the benefits of annotation selection [132] may improve prediction further. Incorporating other types of external information such as transcriptomics through other integrative analysis frameworks such as the transcriptome-wide association study may have added benefits. Combining PGS scores from different methods

and across multiple GWAS sources and distinct populations, in a principled way, such as through bagging or boosting, may ensure robust prediction performance. Incorporating rare genetic variants, especially those with high penetrance, modeling allele frequency and LD-dependent effect size distributions, and accounting for gene–gene interactions and gene–environment interactions, may all improve prediction. Finally, recent studies have suggested that some fraction of the constructed PGS from particular PGS methods may be correlated with and accounted for by non-genetic risk factors [133]. Thus, investigating the benefits of including the constructed PGS on top of the existing non-genetic risk factors used in the baseline risk model for individual disease or all-cause mortality is especially important for assessing the practical performance of PGS methods and the clinical impact of PGS [133,134].

Acknowledgements

This study was supported by National Institutes of Health (NIH) grant R01HG009124 and National Science Foundation (NSF) grant DMS1712933.

Declaration of interests

No interests are declared.

Supplemental information

Supplemental information associated with this article can be found, in the online version, at <https://doi.org/10.1016/j.tig.2021.06.004>

References

- Andersson, L. and Georges, M. (2004) Domestic-animal genomics: deciphering the genetics of complex traits. *Nat. Rev. Genet.* 5, 202–212
- Frazer, K.A. *et al.* (2009) Human genetic variation and its contribution to complex traits. *Nat. Rev. Genet.* 10, 241–251
- McCarthy, M.I. *et al.* (2008) Genome-wide association studies for complex traits: consensus, uncertainty and challenges. *Nat. Rev. Genet.* 9, 356–369
- Martincorena, I. and Campbell, P.J. (2015) Somatic mutation in cancer and normal cells. *Science* 349, 1483–1489
- Nielsen, R. *et al.* (2011) Genotype and SNP calling from next-generation sequencing data. *Nat. Rev. Genet.* 12, 443–451
- Howie, B. *et al.* (2011) Genotype imputation with thousands of genomes. *G3 (Bethesda)* 1, 457–470
- Das, S. *et al.* (2016) Next-generation genotype imputation service and methods. *Nat. Genet.* 48, 1284–1287
- Wray, N.R. *et al.* (2013) Pitfalls of predicting complex traits from SNPs. *Nat. Rev. Genet.* 14, 507–515
- Robinson, M.R. *et al.* (2014) Explaining additional genetic variation in complex traits. *Trends Genet.* 30, 124–132
- Mavaddat, N. *et al.* (2019) Polygenic risk scores for prediction of breast cancer and breast cancer subtypes. *Am. J. Hum. Genet.* 104, 21–34
- Lambert, S.A. *et al.* (2019) Towards clinical utility of polygenic risk scores. *Hum. Mol. Genet.* 28, R133–R142
- Chasioti, D. *et al.* (2019) Progress in polygenic composite scores in Alzheimer's and other complex diseases. *Trends Genet.* 35, 371–382
- Gibson, G. (2019) On the utilization of polygenic risk scores for therapeutic targeting. *PLoS Genet.* 15, e1008060
- So, H.C. *et al.* (2011) Risk prediction of complex diseases from family history and known susceptibility loci, with applications for cancer screening. *Am. J. Hum. Genet.* 88, 548–565
- Visscher, P.M. *et al.* (2012) Five years of GWAS discovery. *Am. J. Hum. Genet.* 90, 7–24
- Lewis, C.M. and Vassos, E. (2020) Polygenic risk scores: from research tools to clinical instruments. *Genome Med.* 12, 44
- Ibanez, L. *et al.* (2019) Polygenic risk scores in neurodegenerative diseases: a review. *Curr. Genet. Med. Rep.* 7, 22–29
- Visscher, P.M. *et al.* (2017) 10 Years of GWAS discovery: biology, function, and translation. *Am. J. Hum. Genet.* 101, 5–22
- Loos, R.J.F. (2020) 15 years of genome-wide association studies and no signs of slowing down. *Nat. Commun.* 11, 5900
- Yang, J. *et al.* (2010) Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* 42, 565–569
- Makowsky, R. *et al.* (2011) Beyond missing heritability: prediction of complex traits. *PLoS Genet.* 7, e1002051
- Hu, Y. *et al.* (2019) A statistical framework for cross-tissue transcriptome-wide association analysis. *Nat. Genet.* 51, 568–576
- Shao, H. *et al.* (2008) Genetic architecture of complex traits: large phenotypic effects and pervasive epistasis. *Proc. Natl. Acad. Sci. U. S. A.* 105, 19910–19914
- International Schizophrenia Consortium *et al.* (2009) Common polygenic variation contributes to risk of schizophrenia and bipolar disorder. *Nature* 460, 748–752
- Hayes, B.J. *et al.* (2010) Genetic architecture of complex traits and accuracy of genomic prediction: coat colour, milk-fat percentage, and type in Holstein cattle as contrasting model traits. *PLoS Genet.* 6, e1001139
- Boyle, E.A. *et al.* (2017) An expanded view of complex traits: from polygenic to omnigenic. *Cell* 169, 1177–1186
- Manolio, T.A. *et al.* (2009) Finding the missing heritability of complex diseases. *Nature* 461, 747–753
- Genin, E. (2020) Missing heritability of complex diseases: case solved? *Hum. Genet.* 139, 103–113
- Visscher, P.M. *et al.* (2008) Heritability in the genomics era – concepts and misconceptions. *Nat. Rev. Genet.* 9, 255–266
- Zhu, H. and Zhou, X. (2020) Statistical methods for SNP heritability estimation and partition: a review. *Comput. Struct. Biotechnol. J.* 18, 1557–1568
- Vilhjalmsson, B.J. *et al.* (2015) Modeling linkage disequilibrium increases accuracy of polygenic risk scores. *Am. J. Hum. Genet.* 97, 576–592
- Wray, N.R. *et al.* (2019) Complex trait prediction from genome data: contrasting EBV in livestock to PRS in humans: genomic prediction. *Genetics* 211, 1131–1141

33. Meuwissen, T.H.E. *et al.* (2001) Prediction of total genetic value using genome-wide dense marker maps. *Genetics* 157, 1819–1829
34. de Koning, D.J. (2016) Meuwissen *et al.* on genomic selection. *Genetics* 203, 5–7
35. Wray, N.R. *et al.* (2007) Prediction of individual genetic risk to disease from genome-wide association studies. *Genome Res.* 17, 1520–1528
36. Bogdan, R. *et al.* (2018) Polygenic risk scores in clinical psychology: bridging genomic risk to individual differences. *Annu. Rev. Clin. Psychol.* 14, 119–157
37. Torkamani, A. *et al.* (2018) The personal and clinical utility of polygenic risk scores. *Nat. Rev. Genet.* 19, 581–590
38. Jostins, L. and Barrett, J.C. (2011) Genetic risk prediction in complex disease. *Hum. Mol. Genet.* 20, R182–R188
39. Abraham, G. and Inouye, M. (2015) Genomic risk prediction of complex human disease and its clinical application. *Curr. Opin. Genet. Dev.* 33, 10–16
40. Chatterjee, N. *et al.* (2016) Developing and evaluating polygenic risk prediction models for stratified disease prevention. *Nat. Rev. Genet.* 17, 392–406
41. Dudbridge, F. *et al.* (2018) Predictive accuracy of combined genetic and environmental risk scores. *Genet. Epidemiol.* 42, 4–19
42. Shen, X. *et al.* (2020) A phenome-wide association and Mendelian randomisation study of polygenic risk for depression in UK Biobank. *Nat. Commun.* 11, 2301
43. Fritsche, L.G. *et al.* (2020) Cancer PRSweb: an online repository with polygenic risk scores for major cancer traits and their evaluation in two independent biobanks. *Am. J. Hum. Genet.* 107, 815–836
44. Richardson, T.G. *et al.* (2019) An atlas of polygenic risk score associations to highlight putative causal relationships across the human genome. *eLife* 8, e43657
45. Choi, S.W. *et al.* (2020) Tutorial: a guide to performing polygenic risk score analyses. *Nat. Protoc.* 15, 2759–2772
46. Guan, Y. and Stephens, M. (2011) Bayesian variable selection regression for genome-wide association studies and other large-scale problems. *Ann. Appl. Stat.* 5, 1780–1815
47. Zhou, X. and Stephens, M. (2012) Genome-wide efficient mixed-model analysis for association studies. *Nat. Genet.* 44, 821–824
48. Habier, D. *et al.* (2011) Extension of the Bayesian alphabet for genomic selection. *BMC Bioinformatics* 12, 186
49. Newcombe, P.J. *et al.* (2019) A flexible and parallelizable approach to genome-wide polygenic risk scores. *Genet. Epidemiol.* 43, 730–741
50. Polson, N.G. and Scott, J.G. (2010) Shrink globally, act locally: sparse Bayesian regularization and prediction. *Bayesian Stat.* 9, 105
51. Moser, G. *et al.* (2015) Simultaneous discovery, estimation and prediction analysis of complex traits using a Bayesian mixture model. *PLoS Genet.* 11, e1004969
52. Lloyd-Jones, L.R. *et al.* (2019) Improved polygenic prediction by Bayesian multiple regression on summary statistics. *Nat. Commun.* 10, 5086
53. Euesden, J. *et al.* (2015) PRSice: polygenic risk score software. *Bioinformatics* 31, 1466–1468
54. Privé, F. *et al.* (2019) Making the most of clumping and thresholding for polygenic scores. *Am. J. Hum. Genet.* 105, 1213–1221
55. Robinson, M.R. *et al.* (2017) Genetic evidence of assortative mating in humans. *Nat. Hum. Behav.* 2, 948–954
56. Yang, S. and Zhou, X. (2020) Accurate and scalable construction of polygenic scores in large biobank data sets. *Am. J. Hum. Genet.* 106, 679–693
57. Verbyla, K.L. *et al.* (2009) Accuracy of genomic selection using stochastic search variable selection in Australian Holstein Friesian dairy cattle. *Genet. Res. (Camb.)* 91, 307–311
58. Hoggart, C.J. *et al.* (2008) Simultaneous analysis of all SNPs in genome-wide and re-sequencing association studies. *PLoS Genet.* 4, e1000130
59. Ge, T. *et al.* (2019) Polygenic prediction via Bayesian regression and continuous shrinkage priors. *Nat. Commun.* 10, 1776
60. Tibshirani, R. (1996) Regression shrinkage and selection via the Lasso. *J. R. Stat. Soc. Ser. B Stat. Method* 58, 267–288
61. de los Campos, G. *et al.* (2009) Predicting quantitative traits with regression models for dense molecular markers and pedigree. *Genetics* 182, 375–385
62. Park, T. and Casella, G. (2012) The Bayesian Lasso. *J. Am. Stat. Assoc.* 103, 681–686
63. Polson, N.G. and Scott, J.G. (2013) Data augmentation for non-Gaussian regression models using variance-mean mixtures. *Biometrika* 100, 459–471
64. Mak, T.S.H. *et al.* (2017) Polygenic scores via penalized regression on summary statistics. *Genet. Epidemiol.* 41, 469–480
65. Pattee, J. and Pan, W. (2020) Penalized regression and model selection methods for polygenic scores on summary statistics. *PLoS Comput. Biol.* 16, e1008271
66. Zhou, X. *et al.* (2013) Polygenic modeling with Bayesian sparse linear mixed models. *PLoS Genet.* 9, e1003264
67. Verbyla, K.L. *et al.* (2010) Sensitivity of genomic selection to using different prior distributions. *BMC Proc.* 4, S5
68. Zeng, P. and Zhou, X. (2017) Non-parametric genetic prediction of complex traits with latent Dirichlet process regression models. *Nat. Commun.* 8, 456
69. Zhou, G. and Zhao, H. (2020) A fast and robust Bayesian non-parametric method for prediction of complex traits using summary statistics. *BioRxiv* Published online December 2, 2020. <https://doi.org/10.1101/2020.11.30.405241>
70. Mak, T.S. *et al.* (2016) Local true discovery rate weighted polygenic scores using GWAS summary data. *Behav. Genet.* 46, 573–582
71. So, H.C. and Sham, P.C. (2017) Improving polygenic risk prediction from summary statistics by an empirical Bayes approach. *Sci. Rep.* 7, 41262
72. Robbins, H. (1956) An empirical Bayes approach to statistics. In *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability* (Vol. 1) (Neyman, J., ed.), pp. 157–163, University of California Press
73. Timpson, N.J. *et al.* (2018) Genetic architecture: the shape of the genetic contribution to human traits and disease. *Nat. Rev. Genet.* 19, 110–124
74. International Multiple Sclerosis Genetics Consortium (2010) Evidence for polygenic susceptibility to multiple sclerosis – the shape of things to come. *Am. J. Hum. Genet.* 86, 621–625
75. Zhang, Y. *et al.* (2018) Estimation of complex effect-size distributions using summary-level statistics from genome-wide association studies across 32 complex traits. *Nat. Genet.* 50, 1318–1326
76. Dudbridge, F. (2013) Power and predictive accuracy of polygenic risk scores. *PLoS Genet.* 9, e1003348
77. Chatterjee, N. *et al.* (2013) Projecting the performance of risk prediction based on polygenic analyses of genome-wide association studies. *Nat. Genet.* 45, 400–405
78. Andlauer, T.F.M. and Nothen, M.M. (2020) Polygenic scores for psychiatric disease: from research tool to clinical application. *Med. Genet.* 32, 39–45
79. Jia, Y. and Jannink, J.L. (2012) Multiple-trait genomic selection methods increase genetic value prediction accuracy. *Genetics* 192, 1513–1522
80. Zhu, X. and Stephens, M. (2017) Bayesian large-scale multiple regression with summary statistics from genome-wide association studies. *Ann. Appl. Stat.* 11, 1561–1592
81. Chun, S. *et al.* (2020) Non-parametric polygenic risk prediction via partitioned GWAS summary statistics. *Am. J. Hum. Genet.* 107, 46–59
82. Privé, F. *et al.* (2020) LDpred2: better, faster, stronger. *Bioinformatics* 36, 5424–5431
83. Ni, G. *et al.* (2020) A comprehensive evaluation of polygenic score methods across cohorts in psychiatric disorders. *MedRxiv* Published online May 5, 2020. <https://doi.org/10.1101/2020.09.10.20192310>
84. Carithers, L.J. *et al.* (2015) A novel approach to high-quality postmortem tissue procurement: the GTEx project. *Biopreserv. Biobank.* 13, 311–319
85. Dixon, J.R. *et al.* (2015) Chromatin architecture reorganization during stem cell differentiation. *Nature* 518, 331–336
86. Kellis, M. *et al.* (2014) Defining functional DNA elements in the human genome. *Proc. Natl. Acad. Sci. U. S. A.* 111, 6131–6138

87. 1000 Genomes Project Consortium (2012) An integrated map of genetic variation from 1,092 human genomes. *Nature* 491, 56–65
88. Kircher, M. *et al.* (2014) A general framework for estimating the relative pathogenicity of human genetic variants. *Nat. Genet.* 46, 310–315
89. GTEx Consortium (2015) The genotype-tissue expression (GTEx) pilot analysis: multitissue gene regulation in humans. *Science* 348, 648–660
90. Roadmap Epigenomics Consortium (2015) Integrative analysis of 111 reference human epigenomes. *Nature* 518, 317–330
91. ENCODE Project Consortium (2012) An integrated encyclopedia of DNA elements in the human genome. *Nature* 489, 57–74
92. Pickrell, J.K. (2014) Joint analysis of functional genomic data and genome-wide association studies of 18 human traits. *Am. J. Hum. Genet.* 94, 559–573
93. Gusev, A. *et al.* (2014) Partitioning heritability of regulatory and cell-type-specific variants across 11 common diseases. *Am. J. Hum. Genet.* 95, 535–552
94. Kichaev, G. *et al.* (2014) Integrating functional data to prioritize causal variants in statistical fine-mapping studies. *PLoS Genet.* 10, e1004722
95. Shi, J. *et al.* (2016) Winner's curse correction and variable thresholding improve performance of polygenic risk modeling based on genome-wide association study summary-level data. *PLoS Genet.* 12, e1006493
96. Speed, D. and Balding, D.J. (2014) MultiBLUP: improved SNP-based prediction for complex traits. *Genome Res.* 24, 1550–1557
97. Hu, Y. *et al.* (2017) Leveraging functional annotations in genetic risk prediction for human complex diseases. *PLoS Comput. Biol.* 13, e1005589
98. Marquez-Luna, C. *et al.* (2020) LDpred-funct: incorporating functional priors improves polygenic prediction accuracy in UK Biobank and 23andMe data sets. *BioRxiv* Published online May 13, 2020. <https://doi.org/10.1101/375337>
99. Maier, R.M. *et al.* (2018) Improving genetic prediction by leveraging genetic correlations among human diseases and traits. *Nat. Commun.* 9, 989
100. Maier, R. *et al.* (2015) Joint analysis of psychiatric disorders increases accuracy of risk prediction for schizophrenia, bipolar disorder, and major depressive disorder. *Am. J. Hum. Genet.* 96, 283–294
101. Hu, Y. *et al.* (2017) Joint modeling of genetically correlated diseases and functional annotations increases accuracy of polygenic risk prediction. *PLoS Genet.* 13, e1006836
102. Zhou, X. and Stephens, M. (2014) Efficient multivariate linear mixed model algorithms for genome-wide association studies. *Nat. Methods* 11, 407–409
103. Gao, B. *et al.* (2021) Accurate genetic and environmental covariance estimation with composite likelihood in genome-wide association studies. *PLoS Genet.* 17, e1009293
104. Li, C. *et al.* (2014) Improving genetic risk prediction by leveraging pleiotropy. *Hum. Genet.* 133, 639–650
105. Chung, W. *et al.* (2019) Efficient cross-trait penalized regression increases prediction accuracy in large cohorts using secondary phenotypes. *Nat. Commun.* 10, 569
106. Chen, T.-H. *et al.* (2020) A penalized regression framework for building polygenic risk models based on summary statistics from genome-wide association studies and incorporating external information. *J. Am. Stat. Assoc.* 116, 133–143
107. Badre, A. *et al.* (2021) Deep neural network improves the estimation of polygenic risk scores for breast cancer. *J. Hum. Genet.* 66, 359–369
108. Bellot, P. *et al.* (2018) Can deep learning improve genomic prediction of complex human traits? *Genetics* 210, 809–819
109. Song, S. *et al.* (2020) Leveraging effect size distributions to improve polygenic risk scores derived from summary statistics of genome-wide association studies. *PLoS Comput. Biol.* 16, e1007565
110. Pal Choudhury, P. *et al.* (2020) iCARE: an R package to build, validate and apply absolute risk models. *PLoS One* 15, e0228198
111. Song, L. *et al.* (2019) SummaryAUC: a tool for evaluating the performance of polygenic risk prediction models in validation datasets with only summary level statistics. *Bioinformatics* 35, 4038–4044
112. Duncan, L. *et al.* (2019) Analysis of polygenic risk score usage and performance in diverse human populations. *Nat. Commun.* 10, 3328
113. Martin, A.R. *et al.* (2017) Human demographic history impacts genetic risk prediction across diverse populations. *Am. J. Hum. Genet.* 100, 635–649
114. Martin, A.R. *et al.* (2019) Clinical use of current polygenic risk scores may exacerbate health disparities. *Nat. Genet.* 51, 584–591
115. Marquez-Luna, C. *et al.* (2017) Multiethnic polygenic risk scores improve risk prediction in diverse populations. *Genet. Epidemiol.* 41, 811–823
116. Weissbrod, O. *et al.* (2021) Leveraging fine-mapping and non-European training data to improve trans-ethnic polygenic risk scores. *MedRxiv* Published online January 20, 2021. <https://doi.org/10.1101/2021.01.19.21249483>
117. Ruan, Y. *et al.* (2021) Improving polygenic prediction in ancestrally diverse populations. *MedRxiv* Published online January 2, 2021. <https://doi.org/10.1101/2020.12.27.20248738>
118. Pain, O. *et al.* (2021) Evaluation of polygenic prediction methodology within a reference-standardized framework. *PLoS Genet.* 17, e1009021
119. Kulm, S. *et al.* (2021) Benchmarking the accuracy of polygenic risk scores and their generative methods. *MedRxiv* Published online March 8, 2021. <https://doi.org/10.1101/2020.04.06.20055574>
120. Yang, J. *et al.* (2011) GCTA: a tool for genome-wide complex trait analysis. *Am. J. Hum. Genet.* 88, 76–82
121. Zhou, X. (2017) A unified framework for variance component estimation with summary statistics in genome-wide association studies. *Ann. Appl. Stat.* 11, 2027–2051
122. Lee, S.H. *et al.* (2011) Estimating missing heritability for disease from genome-wide association studies. *Am. J. Hum. Genet.* 88, 294–305
123. Golan, D. *et al.* (2014) Measuring missing heritability: inferring the contribution of common variants. *Proc. Natl. Acad. Sci. U. S. A.* 111, E5272–E5281
124. Lea, A.J. *et al.* (2015) A flexible, efficient binomial mixed model for identifying differential DNA methylation in bisulfite sequencing data. *PLoS Genet.* 11, e1005650
125. Sun, S. *et al.* (2017) Differential expression analysis for RNAseq using Poisson mixed models. *Nucleic Acids Res.* 45, e106
126. Sun, S. *et al.* (2019) Heritability estimation and differential analysis of count data with generalized linear mixed models in genomic sequencing studies. *Bioinformatics* 35, 487–496
127. Bulik-Sullivan, B. *et al.* (2015) An atlas of genetic correlations across human diseases and traits. *Nat. Genet.* 47, 1236–1241
128. Finucane, H.K. *et al.* (2015) Partitioning heritability by functional annotation using genome-wide association summary statistics. *Nat. Genet.* 47, 1228–1235
129. Yang, J. *et al.* (2015) Genetic variance estimation with imputed variants finds negligible missing heritability for human height and body mass index. *Nat. Genet.* 47, 1114–1120
130. Bulik-Sullivan, B.K. *et al.* (2015) LD Score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nat. Genet.* 47, 291–295
131. Speed, D. and Balding, D.J. (2019) SumHer better estimates the SNP heritability of complex traits from summary statistics. *Nat. Genet.* 51, 277–284
132. Zeng, P. *et al.* (2018) Pleiotropic mapping and annotation selection in genome-wide association studies with penalized Gaussian mixture models. *Bioinformatics* 34, 2797–2807
133. Meisner, A. *et al.* (2020) Combined utility of 25 disease and risk factor polygenic risk scores for stratifying risk of all-cause mortality. *Am. J. Hum. Genet.* 107, 418–431
134. Kulm, S. *et al.* (2020) A systematic framework for assessing the clinical impact of polygenic risk scores. *MedRxiv* Published online April 8, 2020. <https://doi.org/10.1101/2020.04.06.20055574>