



Exploring the Link Between Additive Heritability and Prediction Accuracy From a Ridge Regression Perspective

Arthur Frouin^{1*}, Claire Dandine-Roulland¹, Morgane Pierre-Jean¹, Jean-François Deleuze^{1,2}, Christophe Ambroise^{3*†} and Edith Le Floch^{1*†}

OPEN ACCESS

Edited by:

Lide Han,
Vanderbilt University Medical Center,
United States

Reviewed by:

Kui Zhang,
Michigan Technological University,
United States
Xiaowei Wu,
Virginia Tech, United States
Guo-Bo Chen,
Zhejiang Provincial People's Hospital,
China

*Correspondence:

Christophe Ambroise
christophe.ambroise@univ-evry.fr
Arthur Frouin
arthurs.frouin@gmail.com
Edith Le Floch
edith.lefloch@cnrgh.fr

[†]These authors share last authorship

Specialty section:

This article was submitted to
Statistical Genetics and Methodology,
a section of the journal
Frontiers in Genetics

Received: 09 July 2020

Accepted: 29 September 2020

Published: 04 November 2020

Citation:

Frouin A, Dandine-Roulland C,
Pierre-Jean M, Deleuze J-F,
Ambroise C and Le Floch E (2020)
Exploring the Link Between Additive
Heritability and Prediction Accuracy
From a Ridge Regression Perspective.
Front. Genet. 11:581594.
doi: 10.3389/fgene.2020.581594

¹ CNRGH, Institut Jacob, CEA - Université Paris-Saclay, Évry, France, ² Centre d'Etude du Polymorphisme Humain, Fondation Jean Dausset, Paris, France, ³ LaMME, Université Paris-Saclay, CNRS, Université d'Évry val d'Essonne, Évry, France

Genome-Wide Association Studies (GWAS) explain only a small fraction of heritability for most complex human phenotypes. Genomic heritability estimates the variance explained by the SNPs on the whole genome using mixed models and accounts for the many small contributions of SNPs in the explanation of a phenotype. This paper approaches heritability from a machine learning perspective, and examines the close link between mixed models and ridge regression. Our contribution is two-fold. First, we propose estimating genomic heritability using a predictive approach via ridge regression and Generalized Cross Validation (GCV). We show that this is consistent with classical mixed model based estimation. Second, we derive simple formulae that express prediction accuracy as a function of the ratio $\frac{n}{p}$, where n is the population size and p the total number of SNPs. These formulae clearly show that a high heritability does not imply an accurate prediction when $p > n$. Both the estimation of heritability via GCV and the prediction accuracy formulae are validated using simulated data and real data from UK Biobank.

Keywords: heritability, prediction accuracy, ridge regression, mixed model, generalized cross validation, best linear unbiased predictor

1. INTRODUCTION

The old nature vs. nurture debate is about whether a complex human trait is determined by a person's genes or by the environment. It is a longstanding philosophical question that has been reinvestigated in the light of statistical genetics (Feldman and Lewontin, 1975). The concept of heritability was introduced by Fisher (1918) and Wright (1920, 1921) in the context of pedigree data. It has proved highly useful in animal (Meuwissen et al., 2001) and plant genetics (Xu, 2003) for selection purposes because of its association with accurate prediction of a trait from genetic data. In the last decades, Genome-Wide Association Studies (GWAS) have become highly popular for identifying variants associated with complex human traits (Hirschhorn and Daly, 2005). They have recently been used for heritability estimations (Yang et al., 2010). A shortcut is often made between the heritability of a trait and the prediction of this trait. However, heritable complex human traits are often caused by a large number of genetic variants that individually make small contributions to the trait variation, which is often referred to as polygeny. In this context, the relation between heritability and prediction accuracy may not hold (de Vlaming and Groenen, 2015).

The goal of this paper is to establish a clear relation between prediction accuracy and heritability, especially when the number of genetic markers is much higher than the population size, which is typically the case in GWAS. Based on the linear model, statistical analyses of SNP data address very different and sometimes unrelated questions. The most commonly performed analyses tend to be association studies, where multiple hypothesis testing makes it possible to test the link between any SNP and a phenotype of interest. In genomic selection, markers are selected to predict a phenotype with a view to selecting an individual in a breeding population. Association studies and genomic selection may identify different sets of markers, since even weak associations might be of interest for prediction purposes, while not all strongly associated markers are necessarily useful, because of redundancy through linkage disequilibrium. Genomic heritability allows quantifying the amount of genomic information relative to a given phenotype via mixed model parameter estimation. The prediction of the phenotype using all genomic information via the mixed model is a closely related but different problem.

We approach the problem of heritability estimation from a machine learning perspective. This is not a classical approach in genetics, where inferential statistics is the usual preferred tool. In this context, heritability is considered as a parameter to be inferred from a small sample of the population. The machine learning approach places the emphasis on prediction accuracy. It makes a clear distinction between performance on training samples and performance on testing samples, whereas inferential statistics focuses on parameter estimation on a single dataset.

1.1. Classical Approach via Mixed Models

Heritability is defined as the proportion of phenotypic variance due to genetic factors. A quantitative definition of heritability requires a statistical model. The model commonly adopted is a simple three-term model without gene-environment interaction (Henderson, 1975):

$$\mathbf{y} = \mathbf{g} + \mathbf{f} + \mathbf{e},$$

where $\mathbf{y} \in \mathbb{R}^n$ is a quantitative phenotype vector describing n individuals, $\mathbf{f} \in \mathbb{R}^n$ is a non-genetic covariate term, $\mathbf{g} \in \mathbb{R}^n$ is a genetic term and $\mathbf{e} \in \mathbb{R}^n$ an environmental residual term. The term \mathbf{g} will depend on the diploid genotype matrix $\mathbf{M} \in \mathcal{M}_{n,p}(\mathbb{R})$ of the p causal variants.

There are two definitions of heritability in common use: first, there is H^2 , heritability in the broad sense, measuring the overall contribution of the genome; and second, there is h^2 , heritability in the narrow sense (also known as additive heritability), defined as the proportion of phenotypic variance explained by the additive effects of variants.

The quantification of narrow-sense heritability goes back to family studies by Fisher (1918), who introduced the above model with the additional hypothesis that \mathbf{g} is the sum of independent genetic terms, and with \mathbf{e} assumed to be normal. This heritability in the narrow sense is a function of the correlation between the phenotypes of relatives.

Although Fisher's original model makes use of pedigrees for parameter estimation, some geneticists have proposed using the same model with genetic data from unrelated individuals (Yang et al., 2011a).

1.1.1. Polygenic Model

In this paper, we focus on the version of the additive polygenic model with a Gaussian noise where $\mathbf{g} = \mathbf{Z}\mathbf{u}$, $\mathbf{f} = \mathbf{X}\boldsymbol{\beta}$, with $\mathbf{Z} \in \mathcal{M}_{n,p}(\mathbb{R})$ a standardized (by columns) version of \mathbf{M} , $\mathbf{u} \in \mathbb{R}^p$ a vector of genetic effects, $\mathbf{X} \in \mathcal{M}_{n,r}(\mathbb{R})$ a matrix of covariates, $\boldsymbol{\beta} \in \mathbb{R}^r$ a vector of covariate effects, μ an intercept and $\mathbf{e} \sim \mathcal{N}(\mathbf{0}_n, \sigma^2 \mathbf{I}_n)$ a vector of environmental effects.

The model thus becomes

$$\mathbf{y} = \mu \mathbf{1}_n + \mathbf{Z}\mathbf{u} + \mathbf{X}\boldsymbol{\beta} + \mathbf{e}, \quad (1)$$

where $\mathbf{1}_n \in \mathbb{R}^n$ a vector of ones.

1.1.2. Estimation of Heritability From GWAS Results

To estimate heritability in a GWAS context, a first intuitive approach would be to estimate u with a least squares regression to solve problem (1). Unfortunately, this is complicated in practice for three reasons: the causal variants are not usually available among genotyped variants; genotyped variants are in linkage disequilibrium (LD); and the least squares estimate is only defined when $n > p$, which is not often the case in a GWAS (Yang et al., 2010).

One technique for obtaining a solvable problem is to use the classical GWAS approach to determine a subset of variants significantly associated with the phenotype. The additive heritability can then be estimated by summing their effects estimated by simple linear regressions. In practice this estimation tends to greatly underestimate h^2 (Manolio et al., 2009). It only takes into account variants that have passed the significance threshold after correction for multiple comparisons (strong effects) and does not capture the variants that are weakly associated with the phenotype (weak effects).

1.1.3. Estimating Heritability via the Variance of the Effects

Yang et al. (2010) suggest that most of the missing heritability comes from variants with small effects. In order to be able to estimate the information carried by weak effects they assume a linear mixed model where the vector of random genetic effects follows a normal homoscedastic distribution $\mathbf{u} \sim \mathcal{N}(0_p, \tau \mathbf{I}_p)$. They propose estimating the variance components τ and σ^2 , and defining genomic heritability as $h_G^2 = \frac{p\tau}{p\tau + \sigma^2}$. An example of an algorithm for estimating variance components is the Average Information—Restricted Maximum Likelihood (AI-REML) algorithm, implemented in software such as Genome-wide Complex Trait Analysis (GCTA) (Yang et al., 2011a) or *gaston* (Perdry and Dandine-Roulland, 2018). More recent methods that are much faster than REML have also been proposed, such as the modified Haseman-Elston regression (Chen, 2014) or methods based on summary statistics such as the LD-score regression (Bulik-Sullivan et al., 2015) or the MQS (MinQue for Summary statistics) approach (Zhou, 2017).

1.2. A Statistical Learning Approach via Ridge Regression

The linear model is used in statistical genetics for exploring and summarizing the relation between a phenotype and one or more genetic variants, and it is also used in predictive medicine and genomic selection for prediction purposes. When used for prediction, the criterion for assessing performance is the prediction accuracy.

Although least squares linear regression is the baseline method for quantitative phenotype prediction, it has some limitations. As mentioned earlier, the estimator is not defined when the number of descriptive variables p is greater than the number of individuals n . Even when $n > p$, the estimator may be highly variable when the descriptive variables are correlated, which is clearly the case in genetics.

Ridge regression is a penalized version of least squares that can overcome these limitations (Hoerl and Kennard, 1970). Ridge regression is strongly related to the mixed model and is prediction-oriented.

1.2.1. Ridge Regression

The ridge criterion builds on the least squares criterion, adding an extra penalization term. The penalization term is proportional to the ℓ_2 norm of the parameter vector. The proportionality coefficient λ is also called the penalization parameter. The penalty tends to shrink the coefficients of the least squares estimator, but never cancels them out. The degree of shrinkage is controlled by λ : the higher the value of λ , the greater the shrinkage:

$$\hat{u}_R = \arg \min_u \|y - Zu\|_2^2 + \lambda \|u\|_2^2, \quad (2)$$

$$= \left(Z^T Z + \lambda I_p \right)^{-1} Z^T y, \quad (3)$$

$$= Z^T \left(ZZ^T + \lambda I_n \right)^{-1} y. \quad (4)$$

Ridge regression can be seen as a *Bayesian Maximum a Posteriori* estimation of the linear regression parameters considering a Gaussian prior with hyperparameter λ .

The estimator depends on a λ that needs to be chosen. In a machine learning framework, a classical procedure is to choose the λ that minimizes the squared loss over new observations.

The practical effect of the penalty term is to add a constant to the diagonal of the covariance matrix, which makes the matrix non-singular, even in the case where $p > n$. When the descriptive variables are highly correlated, this improves the conditioning of the $Z^T Z$ matrix, while reducing the variance of the estimator.

The existence theorem states that there always exists a value of $\lambda > 0$ such that the Mean Square Error (MSE) of the ridge regression estimator (variance plus the squared bias) is smaller than the MSE of the Maximum Likelihood estimator (Hoerl and Kennard, 1970). This is because there is always an advantageous bias-variance compromise that reduces the variance without greatly increasing the bias.

Ridge regression also allows us to simultaneously estimate all the additive effects of the genetic variants without discarding any, which reflects the idea that all the variants make a small contribution.

1.2.2. Link Between Mixed Model and Ridge Regression

This paper builds on the parallel between BLUPs (Best Linear Unbiased Predictions) derived from the mixed model and ridge regression (Meuwissen et al., 2001). The use of ridge regression in quantitative genetics has already been discussed (De los Campos et al., 2013; de Vlaming and Groenen, 2015). We look at a machine-learning oriented paradigm for estimating the ridge penalty parameter, which provides us with a direct link to heritability. There is an equivalence between maximizing the posterior $p(u|y)$ and minimizing a ridge criterion (Bishop, 2006) under the assumptions that $u \sim \mathcal{N}(0_p, \tau I_p)$ and $e \sim \mathcal{N}(0_n, \sigma^2 I_n)$ (see section 6 in **Supplementary Material** for details). The optimal penalty hyperparameter of the ridge criterion λ can be used to estimate the heritability. It is indeed defined as the ratio of the variance parameters of the mixed model:

$$\arg \max_u p(u|y) = \arg \min_u \|y - Zu\|_2^2 + \lambda \|u\|_2^2 \text{ with } \lambda = \frac{\sigma^2}{\tau}. \quad (5)$$

The relation between λ and h_G^2 (de Vlaming and Groenen, 2015) is thus:

$$h_G^2 = \frac{p}{p + \lambda}; \lambda = p \frac{1 - h_G^2}{h_G^2}. \quad (6)$$

Consequently, the BLUP has a similar formulation to the ridge estimator. Indeed, as shown in the section 6.2 of the article by Robinson et al. (1991), its general definition is:

$$\hat{u}_{BLUP} = \widehat{\mathbb{E}(u|y)} = \hat{W} Z^T \hat{\Sigma}^{-1} (y - X \hat{\beta}), \quad (7)$$

where $u \sim \mathcal{N}(0_p, W)$, $e \sim \mathcal{N}(0_n, E)$ and $\Sigma = ZWZ^T + E$.

When we further assume $\beta = 0_r$, $W = \tau I_p$ and $E = \sigma^2 I_n$, it becomes:

$$\hat{u}_{BLUP} = \tau Z^T (\tau ZZ^T + \sigma^2 I_n)^{-1} y = Z^T (ZZ^T + \frac{\sigma^2}{\tau} I_n)^{-1} y, \quad (8)$$

which is exactly the ridge estimator.

1.2.3. Over-Fitting

Interestingly, ridge regression and the mixed model can be seen as two similar ways to deal with the classical over-fitting issue in machine learning, which is where a learner becomes overspecialized in the dataset used for the estimation of its parameters and is unable to generalize (Bishop, 2006). When $n > p$, estimating the parameters of a fixed-effect linear model via maximum likelihood estimation may lead to over-fitting, when too many variables are considered. A classical way of reducing over-fitting is regularization, and in order to set the value of the regularization parameter there are two commonly adopted approaches: first, the Bayesian approach, and second, the use of additional data.

Mixed Model parameter estimation via maximum likelihood can be seen as a type of self-regularizing approach (see Equation 5). Estimating the variance components of the mixed model may

be interpreted as a kind of empirical Bayes approach, where the ratio of the variances is the regularization parameter that is usually estimated using a single dataset. In contrast to this, in order to properly estimate the ridge regression regularization hyperparameter that gives the best prediction, two datasets are required. If a single dataset were to be used, this would result in an insufficiently regularized (i.e., excessively complex) model offering too high prediction performances on the present dataset but unable to predict new samples well. This over-fitting phenomenon is particularly evident when dimensionality is high.

The fact that the complexity of the ridge model is controlled by its hyperparameter can be intuitively understood when considering extreme situations. When λ tends to infinity, the estimated effect vector (i.e., \hat{u}_R) tends to the null vector. Conversely, when λ tends to zero, the model approaches maximum complexity. One solution for choosing the right complexity is therefore to use both a training set to estimate the effect vector for different values of the hyperparameter and a validation set to choose the hyperparameter value with the best prediction capacity on this independent sample. An alternative solution, when data is sparse, is to use a cross-validation approach to mimic a two-set situation. Finally, it should be noted that the estimation of prediction performance on a validation dataset is still overoptimistic, and consequently a third dataset, known as a test set, is required to assess the real performance of the model.

1.2.4. Prediction Accuracy in Genetics

In genomic selection and in genomic medicine, several authors have been interested in predicting complex traits that show a relatively high heritability using mixed model BLUPs (Speed and Balding, 2014). The literature defined the prediction accuracy as the correlation between the trait and its prediction, which is unusual in machine learning where the expected loss is often preferred. Several approximations of this correlation have been proposed in the literature (Brard and Ricard, 2015), either in a low-dimensional context (where the number of variants is lower than the number of individuals) or in a high-dimensional context.

Daetwyler et al. (2008) derived equations for predicting the accuracy of a genome-wide approach based on simple least-squares regressions for continuous and dichotomous traits. They consider one univariate linear regression per variant (with a fixed effect) and combine them afterwards, which is equivalent to a Polygenic Risk Score (PRS) (Pharoah et al., 2002; Purcell et al., 2009). Goddard (2009) extended this prediction to Genomic BLUP (GBLUP), which used the concept of an effective number of loci. Rabier et al. (2016) proposed an alternative correlation formula conditionally on a given training set. Their formula refines the formula proposed by Daetwyler et al. (2008). Elsen (2017) used a Taylor development to derive the same formula in small dimension.

Using intensive simulation studies, de Vlaming and Groenen (2015) showed a strong link between PRS and ridge regression in terms of prediction accuracy, when the population size is limited. However, with ridge regression, predictive accuracy improves substantially as the sample size increases.

It is important to note a difference in the prediction accuracy of GBLUP when dealing with human populations as

opposed to breeding populations (De los Campos et al., 2013). De los Campos et al. (2013) show that the squared correlation between GBLUP and the phenotype reaches the trait heritability, asymptotically when considering unrelated human subjects. Dandine-Roulland and Perdry (2015) also proposed a theoretical formula of the performance of BLUPs for prediction in the context of human genetics, which is proportional to the number of individuals, to the squared heritability and to the variance of the off-diagonal terms of the Genetic Relatedness Matrix.

Zhao and Zhu (2019) studied cross trait prediction in high dimension. They derive generic formulae for in and out-of sample squared correlation. They link the marginal estimator to the ridge estimator and to GBLUP. Their results are very generic and generalize formulae proposed by Daetwyler et al. (2008).

1.2.5. Outline of the Paper

While some authors have proposed making use of the equivalence between ridge regression and the mixed model for setting the hyperparameter of ridge regression according to the heritability estimated by the mixed model, we propose on the contrary to estimate the optimal ridge hyperparameter using a predictive approach via Generalized Cross Validation. We derive approximations of the squared correlation and of the expected loss, both in high and low dimensions.

Using synthetic data and real data from UK Biobank, we show that our results are consistent with classical mixed model based estimation and that our approximations are valid.

Finally, with reference to the ridge regression estimation of heritability, we discuss how heritability is linked to prediction accuracy in highly polygenic contexts.

2. MATERIALS AND METHODS

2.1. Generalized Cross Validation for Speeding Up Heritability Estimation via Ridge Regression

2.1.1. Generalized Cross Validation

A classical strategy for choosing the ridge regression hyperparameter uses a grid search and k -fold cross validation. Each grid value of the hyperparameter is evaluated by the cross validated error. This approach is time-consuming in high dimension, since each grid value requires k estimations. In the machine learning context, we propose using Generalized Cross Validation (GCV) to speed up the estimation of the hyperparameter λ and thus to estimate the additive heritability h_G^2 using the link described in Equation (6).

The GCV error in Equation (9) (Golub et al., 1978) is an approximation of the Leave-One-Out error (LOO) (see section 2 in **Supplementary Material**). Unlike the classical LOO, GCV does not require n ridge regression estimations (where n is the number of observations) at each grid value, but involves a single run. It thus provides a much faster and convenient alternative for choosing the hyperparameter. We have

$$\text{err}^{GCV} = \frac{\|\mathbf{y} - \hat{\mathbf{y}}(\lambda)\|_2^2}{\left[\frac{1}{n} \text{tr}(\mathbf{I}_n - \mathbf{H}_\lambda)\right]^2}, \quad (9)$$

where $\hat{y}(\lambda) = \mathbf{Z}\hat{u}_R(\lambda) = \mathbf{H}_\lambda \mathbf{y}$ is the prediction of the training set phenotypes using the same training set for the estimation of \hat{u}_R and where \mathbf{H}_λ is defined as (see section 1 in **Supplementary Material** for details):

$$\begin{aligned} \mathbf{H}_\lambda &= \mathbf{Z} \left(\mathbf{Z}^T \mathbf{Z} + \lambda \mathbf{I}_p \right)^{-1} \mathbf{Z}^T \\ &= \mathbf{Z} \mathbf{Z}^T \left(\mathbf{Z} \mathbf{Z}^T + \lambda \mathbf{I}_n \right)^{-1}. \end{aligned}$$

A Singular Value Decomposition (SVD) of the \mathbf{H}_λ can be used advantageously to speed up GCV computation (see section 3 in **Supplementary Material**).

2.1.2. Empirical Centering Can Lead to Issues in the Choice of Penalization Parameter in a High-Dimensional Setting

In high dimensional settings ($p > n$), the use of GCV after empirical centering of the data can lead to a strong bias in the choice of λ and thus in heritability estimation. Let us illustrate the problem with a simple simulation. We simulate a phenotype from synthetic genotype data with a known heritability of $h^2 = 0.25$, $n = 1,000$ individuals, $p = 10,000$ variants and 100% causal variants. The simulation follows the additive polygenic model without intercept or covariates, as described in section 2.3. Before applying GCV, genotypes are standardized in the most naive way: the genotype matrix \mathbf{M} is empirically centered and scaled column-wise, resulting in the matrix \mathbf{Z} . Since we want to mimic an analysis on real data, let us assume that there is a potential intercept in our model (in practice the empirical mean of our simulated phenotype is likely to be non-null):

$$\mathbf{y} = \mu \mathbf{1}_n + \mathbf{Z}u + \mathbf{e}. \tag{10}$$

GCV expects all the variables to be penalized, but penalizing the intercept is not relevant. We therefore consider a natural two-step procedure: first the model's intercept is estimated via the empirical mean of the phenotype $\hat{\mu} = \frac{1}{n} \sum_i y_i$, and, second, GCV is applied on the empirically centered phenotype $\mathbf{y} = \mathbf{y} - \hat{\mu} \mathbf{1}_n$.

Figure 1 shows the GCV error (dotted line). Heritability is strongly overestimated. The GCV error appears to tend toward its minimum as λ approaches 0 (i.e., when h^2 tends to 1).

This is a direct consequence of the empirical standardization of \mathbf{M} and of the phenotype. By centering the columns of \mathbf{M} with the empirical means of those columns, a dependency is introduced, and each line of the resulting standardized genotype matrix \mathbf{Z} becomes a linear combination of all the others. The same phenomenon of dependency can be observed with the phenotype when using empirical standardization. Given the nature of the LOO in general (where each individual is considered successively as a validation set), this kind of standardization introduces a link between the validation set and the training set at each step: the “validation set individual” can be written as a linear combination of the individuals in the training set. In high dimension, this dependency leads to $\text{err}^{LOO} \xrightarrow{\lambda \rightarrow 0} 0$ (see section 4 in **Supplementary Material**), due to over-fitting occurring in the training set.




FIGURE 1 | Example of biased estimation by GCV if $p > n$. We computed the GCV error curve with $n = 1,000$ individuals, $p = 10,000$ causal variants and simulated heritability $h_{sim}^2 = 0.25$. We used a grid of λ corresponding to the grid of heritability $\{0.01, 0.02, \dots, 0.99\}$ using the link described in Equation (6) and computed the GCV error for those λ after empirical standardization of the data (dotted line). The λ that minimizes the GCV error corresponds to the heritability estimation. Here the GCV error tends to its minimum as h^2 tends to 1, and heritability is thus largely over-estimated. The plain line corresponds to the GCV error obtained after correction of this bias by the projection approach (see section 2.1.3), which provides a satisfactory estimation of h^2 . The three vertical lines correspond respectively to the simulated heritability (red line), the heritability estimated using uncorrected GCV (green line) and the heritability estimated using corrected GCV (blue line).

From a GCV perspective, a related consequence of the empirical centering of the genotype data is that the matrix $\mathbf{Z} \mathbf{Z}^T$ has at least one null eigenvalue and an associated constant eigenvector in a high dimensional setting (see section 4 in **Supplementary Material**). This has a direct impact on GCV: using the singular value decomposition of the empirically standardized matrix $\mathbf{Z} = \mathbf{U} \mathbf{D} \mathbf{V}^T$ with $\mathbf{U} \in \mathcal{O}_n(\mathbb{R})$, $\mathbf{V} \in \mathcal{O}_p(\mathbb{R})$ two orthogonal squared matrices spanning, respectively, the lines and columns spaces of \mathbf{Z} while $\mathbf{D} \in \mathcal{M}_{n,p}(\mathbb{R})$ is a rectangular matrix with singular values $\{d_1, \dots, d_n\}$ on the diagonal. In a high dimensional context: $\text{err}^{GCV}(\mathbf{y}, \mathbf{Z}, \lambda) \xrightarrow[\lambda \rightarrow 0]{d_n^2=0} (\mathbf{1}_n^T \mathbf{y})^2$. Performing the “naive” empirical centering of the phenotype results in

$$\text{err}^{GCV}(\mathbf{y} - \hat{\mu} \mathbf{1}_n, \mathbf{Z}, \lambda) \xrightarrow[\lambda \rightarrow 0]{d_n^2=0} (\mathbf{1}_n^T \mathbf{y} - \mathbf{1}_n^T \hat{\mu} \mathbf{1}_n)^2 = 0.$$

The very same problem is observed for a more general model with covariates (see section 4 in **Supplementary Material**).

2.1.3. A First Solution Using Projection

A better solution for dealing with the intercept [and a matrix of covariates $\mathbf{X} \in \mathcal{M}_{n,r}(\mathbb{R})$] in ridge regression is to use a projection matrix as a contrast and to work on the orthogonal of the space spanned by the intercept (and the covariates).

Contrast matrices are a commonly used approach in the field of mixed models for REstricted Maximum Likelihood computations (REML) (Patterson and Thompson, 1971). REML provides maximum likelihood estimation once fixed effects are taken into account. Contrast matrices are used to “remove” fixed effects from the likelihood formula. If we are only interested in the estimation of the component of variance, we do not even need to make this contrast matrix explicit: any semi-orthogonal matrix $\mathbf{C} \in \mathcal{M}_{n-r-1,n}(\mathbb{R})$ such that $\mathbf{C}\mathbf{C}^T = \mathbf{I}_{n-r-1}$ and $\mathbf{C} \times (\mu \mathbf{1}_n + \mathbf{X}\beta) = \mathbf{0}_{n-r-1}$ provides a solution (see section 6 in **Supplementary Material** for details). In a ridge regression context, an explicit expression of \hat{u} is needed for choosing the optimal complexity. An explicit form for \mathbf{C} is therefore necessary.

In the presence of covariates, a QR decomposition can be used to obtain an explicit form for \mathbf{C} (see section 5 in **Supplementary Material** for details). In the special case of an intercept without covariates, there is a convenient choice of \mathbf{C} . Since the eigenvector of $\mathbf{Z}\mathbf{Z}^T$ associated with the final null eigenvalue is constant, $\mathbf{C} = [\mathbf{U}_1, \dots, \mathbf{U}_{n-1}]^T \in \mathcal{M}_{n-1,n}(\mathbb{R})$ is a contrast matrix adapted for our problem. Additionally, by considering $\mathbf{C}\mathbf{Z}$ instead of \mathbf{Z} , we have $\mathbf{C}\mathbf{Z} = \mathbf{D}_{-n}\mathbf{V}^T \rightarrow \mathbf{C}\mathbf{Z}\mathbf{Z}^T\mathbf{C}^T = \mathbf{D}_{-n}\mathbf{D}_{-n}^T$ with \mathbf{D}_{-n} the matrix \mathbf{D} deprived of row n . This choice of contrast matrix thus simplifies the GCV formula and allows extremely fast computation.

2.1.4. A Second Solution Using 2 Data Sets

Dependency between individuals can be a problem when we use the same data for the standardization (including the estimation of potential covariate effects) and for the estimation of the genetic effects. This can be overcome by partitioning our data. Splitting our data into a standardization set and a training set, we will first use the standardization set to estimate the mean and the standard deviation of each variant, the intercept, and the potential covariate effects. Those estimators will then be used to standardize the training set on which GCV can then be applied.

This method has two main drawbacks. The first is that the estimation of the non-penalized effects is done independently of the estimation of the genetic effects, even though in practice we do not expect covariates to be highly correlated with variants. The other drawback is that it reduces the number of individuals for the heritability estimation (which is very sensitive to the number of individuals). This approach therefore requires a larger sample than when using projection.

2.2. Prediction vs. Heritability in the Context of Small Additive Effects

Ridge regression helps to highlight the link between heritability and prediction accuracy. What is the relation between the two concepts? Is prediction accuracy an increasing function of heritability?

In a machine learning setting, we have training and testing sets. The classical bias-variance trade-off formulation considers

the expectation of the loss over both the training set and the test individual phenotype. It breaks down the prediction error into three terms commonly called variance, bias, and irreducible error. In this paper we do consider the genotypes of the training set as fixed and the genotype of a test individual as random, and somewhat abusively continue to employ the terms variance, bias, and irreducible error:

$$\begin{aligned} \mathbb{E}_{\mathbf{y}_{tr}, \mathbf{y}_{te}, z_{te}} [(y_{te} - \hat{y}_{te})^2] &= \mathbb{E}_{z_{te}} [\mathbb{E}_{\mathbf{y}_{tr}, \mathbf{y}_{te} | z_{te}} [(y_{te} - \hat{y}_{te})^2]] \\ &= \mathbb{E}_{z_{te}} [\text{var}(y_{te} | z_{te}) + \text{var}(\hat{y}_{te} | z_{te})] + \\ &\quad \mathbb{E}_{z_{te}} [(\mathbb{E}_{\mathbf{y}_{tr} | z_{te}} [\hat{y}_{te}] - \mathbb{E}_{\mathbf{y}_{te} | z_{te}} [y_{te}])^2]. \end{aligned}$$

where the index tr refers to the training set, while te refers to the test set.

Assuming a training set genotype matrix $\mathbf{Z} \in \mathcal{M}_{n,p}(\mathbb{R})$ (without index tr to lighten notations) whose columns have zero mean and unit variances, we denote $\mathbf{K}_\lambda = (\mathbf{Z}^T\mathbf{Z} + \lambda\mathbf{I}_p)^{-1}\mathbf{Z}^T$. Assuming the independence of the variants $\mathbb{E}_{z_{te}} [z_{te}] = 0_p$ and $\text{var}(z_{te}) = \mathbf{I}_p$, irreducible error, variance, and bias become:

$$\begin{aligned} \mathbb{E}_{z_{te}} [\text{var}(y_{te} | z_{te})] &= \sigma^2 \\ \mathbb{E}_{z_{te}} [\text{var}(\hat{y}_{te} | z_{te})] &= \sigma^2 \text{tr}(\mathbf{K}_\lambda \mathbf{K}_\lambda^T) \\ \mathbb{E}_{z_{te}} [(\mathbb{E}_{\mathbf{y}_{tr} | z_{te}} [\hat{y}_{te}] - \mathbb{E}_{\mathbf{y}_{te} | z_{te}} [y_{te}])^2] &= u^T (\mathbf{K}_\lambda \mathbf{Z} - \mathbf{I}_p)^2 u. \end{aligned}$$

where u is the vector of the ridge parameters.

Since individuals are assumed to be unrelated, the covariance matrix of the individuals is diagonal. The covariance matrix of the variants is also diagonal, since variants are assumed independent. Assuming scaled data, $\mathbf{Z}\mathbf{Z}^T$ and $\mathbf{Z}^T\mathbf{Z}$ are the empirical estimations of covariance matrices of respectively the individuals and the variants (up to a p or n scaling factor). Two separate situations can be distinguished according to the n/p ratio. In the high-dimensional case where $p > n$, the matrix $\mathbf{Z}\mathbf{Z}^T$ estimates well the individuals' covariance matrix up to a factor p . Where $n > p$, on the other hand, $\mathbf{Z}^T\mathbf{Z}$ estimates well the covariance matrix of variants up to a factor n . Eventually, $\mathbf{Z}\mathbf{Z}^T \simeq p\mathbf{I}_n$ when $n < p$ and $\mathbf{Z}^T\mathbf{Z} \simeq n\mathbf{I}_p$ when $n > p$.


Assuming further that

- $\forall i \in \llbracket 1, n \rrbracket \text{ var}(y_i) = 1$, we then have $\sigma^2 = 1 - h^2$,
- heritability is equally distributed among normalized variants i.e., $\forall j \in \llbracket 1, p \rrbracket \text{ var}(u_j) = \frac{h^2}{p}$ (which is indeed the mixed model hypothesis),
- $u^T u \simeq p \times \frac{h^2}{p}$ and $(\mathbf{Z}u)^T (\mathbf{Z}u) \simeq nh^2$,

the expected prediction error can be stated more simply, according to the $\frac{n}{p}$ ratio (see section 8 in **Supplementary Material** for details):

$$\mathbb{E}_{\mathbf{y}_{tr}, \mathbf{y}_{te}, z_{te}} [(y_{te} - \hat{y}_{te})^2] \simeq \begin{cases} 1 - \frac{n}{p}(h^2)^2, & \text{if } p \geq n \\ (1 - h^2) \frac{1 + \frac{n}{p}h^2}{1 + h^2(\frac{n}{p} - 1)}, & \text{otherwise.} \end{cases} \quad (11)$$

When considering the theoretical quadratic error with respect to the log ratio of the number of individuals over the number



of variants in the training set (Figure 2), as expected we have a decreasing function. This means that the larger the number of individuals in the training sample, the smaller the error. We also observe that the higher the heritability, the smaller the error. Both of these things are intuitive, and as a consequence the error tends toward the irreducible error when n becomes much larger than p . What is more surprising is that the prediction error is close to the maximum, whatever the heritability, when n is much smaller than p . Paradoxically, even with the highest possible heritability, if the number of variants is too large in relation to the number of individuals, no prediction is possible. This can be explained by the fact that the penalization plays a very important role in that case and thus strongly increases the bias, while reducing the variance. The squared bias and the variance with respect to the log ratio of the number of individuals over the number of variants in the training set are shown in Supplementary Figures 3, 4. The irreducible error is only a function of heritability and is not affected by the dimension of the training set.

Similarly, the prediction error can be computed on the training set instead of on the test set. Using the same assumptions as before, the expected prediction error on the training set can be approximated by:

$$\mathbb{E}_{\mathbf{y}_{tr}} \left[\frac{1}{n} (\mathbf{y}_{tr} - \hat{\mathbf{y}}_{tr})^T (\mathbf{y}_{tr} - \hat{\mathbf{y}}_{tr}) \right] \approx \begin{cases} (1 - h^2)^2 & \text{if } p > n, \\ 1 - 2 \frac{n}{n+\lambda} \left(\frac{p}{n} (1 - h^2) + h^2 \right) + \left(\frac{n}{n+\lambda} \right)^2 \left(\frac{p}{n} (1 - h^2) + h^2 \right)^2 & \text{otherwise.} \end{cases}$$

A graph similar to Figure 2 for this expected error can be found in Supplementary Figure 5. Interestingly, when $p > n$, the error on the training set does not depend on the n/p ratio. When n becomes greater than p , it increases and tends toward the irreducible error $1 - h^2$ when $n \gg p$. As shown in Figure 2, the error on the test set is always higher than the irreducible error and thus higher than the error on the training set, which is a sign of over-fitting. However, the difference between the error on the test set and the error on the training set is a decreasing function of the n/p ratio, which is linear when $p > n$ and tends toward zero when $n \gg p$.

Another popular way of looking at the predictive accuracy is to consider the squared correlation between y_{te} and \hat{y}_{te} (Goddard, 2009; Daetwyler et al., 2010):

$$\text{corr}^2(y_{te}, \hat{y}_{te}) = \frac{\text{cov}^2(y_{te}, \hat{y}_{te})}{\text{var}[y_{te}] \text{var}[\hat{y}_{te}]}$$

Although correlation and prediction error both provide information about the prediction accuracy, correlation may have an interpretation that is intuitive, but it does not take the scale of the prediction into account. From a predictive point of view, this is clearly a disadvantage. Considering y_{te} , z_{te} , and y_{tr} to be random, and using the same assumptions that were made in relation to prediction error, the three terms of the squared correlation become:

$$\begin{aligned} \text{cov}^2(y_{te}, \hat{y}_{te}) &= (u^T \mathbf{K}_\lambda \mathbf{Z}_{tr} u)^2, \\ \text{var}[\hat{y}_{te}] &= \text{tr}(\mathbf{K}_\lambda^T \mathbf{K}_\lambda \times \sigma^2 \mathbf{I}_n) + (\mathbf{Z}_{tr} u)^T \mathbf{K}_\lambda^T \mathbf{K}_\lambda (\mathbf{Z}_{tr} u), \\ \text{var}[y_{te}] &= 1. \end{aligned}$$


Like in the case of prediction error, replacing $\mathbf{Z}\mathbf{Z}^T$ or $\mathbf{Z}^T\mathbf{Z}$ by their expectations, the squared correlation simplifies to:

$$\text{corr}^2(y_{te}, \hat{y}_{te}) \simeq \begin{cases} \frac{n}{p} (h^2)^2 & \text{if } n < p, \\ \frac{(h^2)^2}{\frac{p}{n} (1 - h^2) + h^2} & \text{otherwise.} \end{cases} \quad (12)$$

When considering this theoretical squared correlation with respect to the log ratio of the number of individuals over the number of variants in the training set (Figure 3), we have, as expected, an increasing function. Similarly, the higher the heritability, the higher the squared correlation. We also observe that when $n \gg p$, the squared correlation tends toward the simulated heritability. Conversely, when $p \gg n$, it is close to zero whatever the heritability.

2.3. Simulations and Real Data

Since narrow-sense heritability is a quantity that relates to a model, we will first illustrate our contributions via simulations where the true model is known. We perform two different types of simulation: fully synthetic simulations where both genotypes and phenotypes are drawn from statistical distributions, and semi-synthetic simulations where UK Biobank genotypes are used to simulate phenotypes. We also illustrate our contributions using height and body mass index (BMI) from the UK Biobank dataset.



We first assess the performance of GCV for heritability estimation and then look at the accuracy of the prediction when the ratio of the number of individuals to the number of variants varies in the training set.

2.3.1. UK Biobank Dataset

The present analyses were conducted under UK Biobank data application number 45,408. The UK Biobank dataset consists of $\simeq 784$ K autosomal SNPs describing $\simeq 488$ K individuals. We applied relatively stringent quality control and minor allele frequency filters to the dataset (callrate for individuals and variants > 0.99 ; p -values of Hardy-Weinberg equilibrium test $> 1e-7$; Minor Allele Frequency > 0.01), leading to 473,054 and 417,106 remaining individuals and SNPs, respectively.

Two phenotypes were considered in our analyses: height (standing) and BMI. In order to have a homogeneous population for the analysis of these real phenotypes, we retained only those individuals who had reported their ethnicity as white British and whose Principal Component Analysis (PCA) results obtained by UK Biobank were consistent with their self-declared ethnicity. In addition, each time we subsampled individuals we removed related individuals [one individual in all pairs with a Genetic Relatedness Matrix (GRM) coefficient > 0.025 was removed], as in Yang et al. (2011b) in order to avoid confusion between shared genetic factors and shared environmental factors. Several covariates were also considered in the analysis of these phenotypes: the sex, the year of birth, the recruitment center,

TABLE 1 | Table of the parameters sets of the simulations.

Parameters	Levels
n/p	Simulation: 1,000/10,000; 5,000/10,000; 10,000/500,000 Data-based: 1,000/10,000; 5,000/10,000; 10,000/417,106
f_c	0.1; 0.5; 1
h_{sim}^2	{0.1, ..., 0.9}

n/p : the ratio of the dimensions of the genotype matrix. f_c : proportion of causal variants. h_{sim}^2 : simulated heritability.

the genotyping array, and the first 10 principal components computed by UK Biobank.

2.3.2. Synthetic Genotype Data

The synthetic genotype matrices are simulated as in Golan et al. (2014) and de Vlaming and Groenen (2015). This corresponds to a scenario with independent loci or perfect linkage equilibrium. To simulate synthetic genotypes for p variants, we first set a vector of variant frequencies $f \in \mathbb{R}^p$, with these frequencies independently following a uniform distribution $\mathcal{U}([0.05, 0.5])$. Individual genotypes are then drawn from binomial distributions with proportions f , to form the genotype matrix \mathbf{M} . A matrix of standardized genotypes \mathbf{Z}^* can be obtained by standardizing \mathbf{M} with the true variant frequencies f .

2.3.3. Simulations to Assess Heritability Estimation Using GCV

We consider both synthetic and real genetic data, and simulate associated phenotypes.

In the two simulation scenarios we investigate the influence on heritability estimation of the following three parameters: the shape of the genotype matrix in the training set (the ratio between n the number of individuals and p the number of variants), the fraction of variants with causal effects f_c , and the true heritability h_{sim}^2 . The tested levels of these quantities are shown in **Table 1**.

For each simulation scenario and for a given a set of parameters (n , p , f_c , h_{sim}^2), the simulation of the phenotype starts with a matrix of standardized genotypes (either a synthetic genotype matrix \mathbf{Z}^* standardized with the true allele frequencies, as described in section 2.3.2, or a matrix of empirically standardized genotypes \mathbf{Z} obtained from UK Biobank data). To create the vector of genotype effects u , $p \times f_c$ causal SNPs are randomly sampled and their effects are sampled from a multivariate normal distribution with zero mean and a covariance matrix $\tau \mathbf{I}_{p \times f_c}$ (where $\tau = \frac{h_{sim}^2}{p \times f_c}$), while the remaining $p \times (1 - f_c)$ effects are set to 0. The vector of environmental effects e is sampled from a multivariate normal distribution with zero mean and a covariance matrix $\sigma^2 \mathbf{I}_n$, where $\sigma^2 = 1 - h_{sim}^2$. The phenotypes are then generated as $\mathbf{y} = \mathbf{Z}^* u + e$ and $\mathbf{y} = \mathbf{Z} u + e$, for the fully synthetic scenario and the semi-synthetic scenario, respectively. A standardization set of 1,000 individuals (that will be used for the GCV approach based on two datasets) is also generated for each scenario in the same way.

Applying GCV to large-scale matrices can be extremely time-consuming, since it requires the computation of the GRM associated with \mathbf{Z}^* or \mathbf{Z} and the eigen decomposition of

the GRM. For this reason we employed the same strategy as de Vlaming and Groenen (2015) in order to speed up both simulations and analyses by making it possible to test more than one combination of simulation parameters. We simulated an ($n_{max} = 10,000 \times p_{max} = 500,000$) genotype matrix for the training set in the fully synthetic scenario and used this simulated matrix for all the $9 \times 3 \times 3 = 81$ ($h_{sim}^2 \times f_c \times n/p$) parameter combinations. Similarly, we sampled $n_{max} = 10,000$ individuals from the UK Biobank dataset to obtain an ($n_{max} = 10,000 \times p_{max} = 417,106$) genotype matrix for the training set in the semi-synthetic scenario. Smaller matrices were then created from a subset of these two large matrices (note that for subsets of the real genotype matrix we took variants in the original order to keep the linkage disequilibrium structure). Consequently, computation of the GRM and its eigen decomposition needed to be performed only once for each n/p ratio considered. The fully synthetic and the semi-synthetic scenarios were each replicated 30 times.

2.3.4. Simulations to Assess Prediction Accuracy

We performed fully synthetic simulations for different ratios $\frac{n}{p}$ in order to study the behavior of the mean prediction error and the correlation between the phenotype and its prediction. We considered a training set of size $n = 1,000$, and a test set of size $n_{te} = 5,000$. The maximum number of variants was set to $p_{max} = 50,000$ and the heritability to $h^2 = 0.6$. We first simulated a global allelic frequency vector $f \sim \mathcal{U}_{p_{max}}(0.05, 0.5)$ and a global vector of genetic effects $u \sim \mathcal{N}\left(p_{max}, \frac{h^2}{p_{max}} \mathbf{I}_{p_{max}}\right)$.

For each subset of variants of size $p < p_{max}$, we selected a vector of genetic effects composed of the p first components of u multiplied by a $\sqrt{\frac{p_{max}}{p}}$ factor assuring a total variance of 1 and $\text{var}(u^p) = \frac{h^2}{p} \mathbf{I}_p$: $u^p = (u_1, \dots, u_p) \times \sqrt{\frac{p_{max}}{p}}$. The genotype matrix \mathbf{M}_{te} was then simulated and its normalized version \mathbf{Z}_{te}^* computed as described in section 2.3.2. The normalization used the first p components of f . The noise vector $\mathbf{e}_{te} \sim \mathcal{N}(\mathbf{0}_{n_{te}}, (1 - h^2)\mathbf{I}_{n_{te}})$ and a vector of phenotypes $\mathbf{y}_{te} = \mathbf{Z}_{te}^* u^p + \mathbf{e}_{te}$ were eventually simulated.

We generated 300 training sets by simulating the normalized genotype matrix, noise, and phenotype using the same process as for the test set. Here, the training set index is denoted as k . A prediction $\hat{\mathbf{y}}_{te,k}$ for the test set was made with each training set using the ridge estimator of u^p obtained with $\lambda = p \frac{1-h^2}{h^2}$, and the following empirical quantities were estimated: $\text{err}_p = \frac{1}{300} \sum_k \frac{1}{n_{te}} \left\| \mathbf{y}_{te,k} - \hat{\mathbf{g}}_p \right\|_2^2$, $\text{bias}_p^2 = \frac{1}{n_{te}} \sum_{i \in [1, n_{te}]} \left(\left[\mathbf{Z}_{te} u^p - \hat{\mathbf{g}}_p \right]_i \right)^2$ and $\text{var}_p = \frac{1}{300} \sum_k \frac{1}{n_{te}} \left\| \hat{\mathbf{y}}_{te,k} - \hat{\mathbf{g}}_p \right\|_2^2$, where $\hat{\mathbf{g}}_p = \left(\frac{1}{300} \sum_{k \in [1, 300]} \left[\hat{\mathbf{y}}_{te,k} \right]_i \right)_{i \in [1, n_{te}]}$. The squared correlation between $\hat{\mathbf{y}}_{te,k}$ and $\mathbf{y}_{te,k}$ was also estimated.

We considered the following numbers of variants:

$$p \in \{50,000; 25,000; 16,667; 12,500; 10,000; 5,000; 3,333; 2,500; 2,000; 1,667; 1,429; 1,250; 1,111; 1,000; 500; 136; 79; 56; 43; 35; 29; 25; 22; 20\}.$$

TABLE 2 | Size (number of individuals) of training, standardization, and test sets for assessing predictive power on real data.

Set	Size
Training	{1,000; 2,000; 5,000; 10,000; 20,000}
Standardization	1,000
Test	1,000

TABLE 3 | Number of repetitions for the evaluation of the predictive power on real data.

Size of the training set	1,000	2,000	5,000	10,000	20,000
Number of repetitions	100	70	50	20	10

2.4. Prediction of Height and BMI Using UK Biobank Data

To experiment on UK Biobank for assessing the prediction accuracy, for each phenotype we considered three sets of data: a training set for the purpose of learning genetic effects, a standardization set for learning non-penalized effects (covariates and intercept), and a test set for assessing predictive power. Pre-treatment filters (as described in section 2.3.1) were systematically applied on the training set. We computed the estimation of genetic effects using the projection-based approach to take into account non-penalized effects, where the penalty parameter was obtained by GCV with the same projection approach:

$$\hat{u}_R = \mathbf{Z}_{tr}^T \mathbf{C}_{tr}^T \left(\mathbf{C}_{tr} \mathbf{Z}_{tr} \mathbf{Z}_{tr}^T \mathbf{C}_{tr}^T + \hat{\lambda}_{GCV} \mathbf{I}_{n-r} \right)^{-1} \mathbf{C}_{tr} \mathbf{y}_{tr}.$$

We then estimated non-penalized effects (here \mathbf{X} contains the intercept):

$$\hat{\beta} = \left(\mathbf{X}_{std}^T \mathbf{X}_{std} \right)^{-1} \mathbf{X}_{std}^T (\mathbf{y}_{std}). \tag{13}$$

Finally, we applied these estimations on the test set:

$$\begin{aligned} \hat{\mathbf{g}}_{te} &= \mathbf{Z}_{te} \hat{u}_R, \\ \hat{\mathbf{f}}_{te} &= \mathbf{X}_{te} \hat{\beta}, \\ \tilde{\mathbf{y}}_{te} &= \mathbf{y}_{te} - \hat{\mathbf{f}}_{te}, \end{aligned}$$

in order to compute the Mean Square Error = $\frac{1}{n_{te}} (\tilde{\mathbf{y}}_{te} - \hat{\mathbf{g}}_{te})^T (\tilde{\mathbf{y}}_{te} - \hat{\mathbf{g}}_{te})$ between the phenotype residuals $\tilde{\mathbf{y}}_{te}$ after removal of non-penalized effects and $\hat{\mathbf{g}}_{te}$.

This procedure was performed for different ratios $\frac{n}{p}$ using different sized subsets of individuals for the training set, while keeping all the variants that passed pre-treatment filters (see **Table 2**).

For each number n of individuals considered in the training set, the sampling of these individuals was repeated several times, as seen in **Table 3**, in order to account for the variance of the estimated genetic effects due to sampling.

3. RESULTS

3.1. Generalized Cross Validation for Heritability Estimation

3.1.1. Simulation Results

For the two simulation scenarios we look at the difference between the estimation of h_g^2 by GCV and the simulated heritability h_{sim}^2 in different configurations of study size n/p , h_{sim}^2 and the fraction of causal variants f_c . Similarly, we look at the difference between the estimation by the classical mixed model approach and the simulated heritability. In our simulations f_c was seen to have no influence, and so only the influence of the remaining parameters is shown in **Figure 4** and f_c is fixed at 10%. For full results, see **Supplementary Figures 1, 2**.

For the fully-simulated scenario, the two GCV approaches give very similar results and appear to provide an unbiased estimator of h^2 . They compare very well with the estimation of heritability by ridge regression with a 10-fold CV. Moreover, the variance of the GCV estimators does not appear higher than the variance of 10-fold CV.

In the case of the semi-synthetic simulations, here too both GCV approaches and the 10-fold CV provide a satisfactory heritability estimation. Our choice of using GCV in place of a classical CV approach for estimating heritability by ridge regression is therefore validated.

For both simulation scenarios we also note that the classical mixed model approach (using the AI-REML method in the `gaston` R package) gives heritability estimations that are very similar to those obtained using the GCV approaches. The value of simulated heritability does not appear to have a strong effect on the quality of the heritability estimation. On the other hand, the ratio n/p seems to have a real impact on estimation variance, with lower ratios leading to lower variances, which initially might appear surprising. One possible explanation for this is that in our simulations n increases as the ratio n/p decreases. Visscher and Goddard (2015) showed that the variance of the heritability is a decreasing function of n , which could explain the observed behavior.

3.1.2. Illustration on UK Biobank

We now compare heritability estimations between the two GCV approaches and the classical mixed model approach for height and BMI, on a training set of 10,000 randomly sampled individuals (the training set being of the same size as for the simulated data). All three approaches take account of covariates and the intercept. The AI-REML approach also uses a projection matrix to deal with covariates. For the GCV approach based on two datasets, a standardization set of 1,000 individuals is also sampled, and for comparison purposes we have chosen to apply this two-set strategy to the classical mixed model approach as well.

Since the true heritability is of course unknown with real data, the sampling of the training and standardization sets is repeated 10 times in order to account for heritability estimation variability. Note that the SNP quality control and MAF filters were repeated at each training set sampling and applied to the standardization set.

Figure 5 shows that for each phenotype the two GCV approaches and the classical mixed model approach (AI-REML) give similar estimations. There is relatively little estimation variability, and any variability observed seems depend more on the individuals sampled for the training set than on the approach used.

3.2. Prediction vs. Heritability in the Context of Small Additive Effects

3.2.1. Prediction From Synthetic Data

As expected, the mean of the test set error follows closely the theoretical curve when the $\log \frac{n}{p}$ varies (**Figure 6**). When $n > p$, the mean of the test set is close to the minimum possible error, which means that the ridge regression provides a reliable prediction on average.

Interestingly, if the mean error behaves as expected by our approximation, the standard deviation of the error may be very large. **Figures 6A,B** show the same mean error with different error bars. **Figure 6A** plots the error bars corresponding to the training set variation: the mean test set error is computed for each training set and the error bars show one standard deviation across the 300 training sets. **Figure 6B** plots the error bars corresponding to the variation of the errors across the test set.

The error bars in **Figure 6B** are much larger than those in **Figure 6A**, which shows that the variation in the prediction error is mostly due to the test individual whose phenotype we wish to predict, and depends little on the training set. This may be explained by the fact that the environmental residual term can be very large for some individuals. For these individuals the phenotype will be predicted with a very large error even when $n \gg p$, that is to say when the genetic term is correctly estimated, irrespective of the training set.

The squared correlation between the phenotype and its prediction, as a function of $\log \frac{n}{p}$, is also in line with our approximation (**Figure 7**). As expected, when $n \gg p$, the squared correlation tends toward the simulated heritability. We compared our approximation with the approximation obtained by Daetwyler et al. (2008) and observed that although Daetwyler's approximation is very similar to ours when $p \gg n$, our simulation results make Daetwyler's approximation appear under-optimistic when $n \gg p$. Finally, we also compared our approximation with that obtained by Rabier et al. (2016), which is the same as ours when $n > p$. However, when $p > n$, Rabier's approximation appears over-optimistic.

3.2.2. Prediction From UK Biobank Data

Let us consider the proposed theoretical approximation of the predictive power of ridge regression with respect to the n/p ratio applied to the UK Biobank data, for height and BMI residuals (after removal of covariate effects and intercept).

The two phenotypes differ considerably as regards heritability: we estimate by the projection-based GCV approach that 73.33% of height is "heritable" whereas only 33.91% of BMI is (on average over the 10 training samples of 20,000 individuals).

These estimated values are close to those currently found in the literature (Ge et al., 2017). It is important to note that




FIGURE 4 | Distribution of $(h^2_{est} - h^2_{sim})$ for different parameter combinations with 30 replications. **(A)** Corresponds to data simulated under the “fully synthetic” procedure, while **(B)** corresponds to the “semi-synthetic simulation” procedure. Each sub-panel corresponds to a different value of n/p . In both scenarios 10% of the variants have causal effects (i.e., $f_c = 0.1$). For each panel, the horizontal axis corresponds to the simulated heritability $h^2_{sim} \in \{0.1, \dots, 0.9\}$ and the vertical axis corresponds to $(h^2_{est} - h^2_{sim})$. Heritability estimations are done with the random effects model using AI-REML and with ridge regression using three approaches for the choice of λ : GCV with a projection correction, GCV with a two-dataset correction and a 10-fold cross-validation (CV 10f).




FIGURE 5 | Heritability estimation of BMI and height using AI-REML and GCV, with the projection-based approach and with the two-set approach. We sub-sampled the original UK Biobank dataset 10 times for replication. The cross corresponds to the mean and the error bar to the mean \pm one standard deviation.

the heritability estimation is strongly dependent on the filters. Variations of up to 20% were observed in the estimations when the filtering procedure setup was slightly modified.

A major difference between UK Biobank data and our simulations designed to check the proposed approximation lies in the strong linkage disequilibrium present in the human genome. Several papers have proposed using the effective number of independent markers to make adjustments in the multiple testing framework (Li et al., 2012), and we likewise propose adjusting our prediction model by taking into account an effective number of SNPs (p_e). We estimate the effective $\frac{n}{p_e}$ ratio for each training set and for each considered n value using the observed mean square errors, the estimated heritability, and the theoretical relation in the case of independent variants $\mathbb{E}_{y_{tr}, y_{te}, z_{te}} [(y_{te} - \hat{y}_{te})^2] = 1 - \frac{n}{p} (h^2)^2$ when $p > n$. We then use a simple linear regression to find the coefficient between these estimated $\frac{n}{p_e}$ ratios and the corresponding real $\frac{n}{p}$ ratios.

Table 4 shows different but close effective numbers of SNPs for the two phenotypes.

We also consider normalizing the test set errors using the mean square error of phenotype residuals (after removing non-penalized effects). Using this error normalization and adjusting the theoretical curve for an effective number of SNPs, we observe

a close fit between the estimated errors on the test set and their theoretical values (**Figure 8**).

4. DISCUSSION

In this work we investigated an alternative computation of genomic heritability based on ridge regression. We proposed a fast, reliable way to estimate the optimal penalization parameter of the ridge via Generalized Cross Validation adapted for high dimension. The genomic heritability estimated from the GCV gives results comparable to mixed model AIREML estimates. It clearly demonstrates that a predictive criterion allows a reliable choice of the penalization parameter and associated heritability, even when the prediction accuracy of ridge regression is low. Moreover, even though our approach does not formally consider Linkage Disequilibrium, simulations showed that it still provides reliable genomic heritability estimates in presence of realistic Linkage Disequilibrium.

We also provided theoretical approximations of the ridge regression prediction accuracy, in terms of both error and correlation between the phenotype and its prediction on new samples. These approximations perform well on synthetic data, in both high and low dimensions. They rely on the assumption that individuals and markers are independent in approximating






TABLE 4 | Effective number of SNPs.

Phenotype	p/p_e
Height	5.01
BMI	3.48

exhibits similar performances to Daetwyler et al. (2008) when $p > n$ and to Rabier et al. (2016) when $p < n$.

Our theoretical approximation of the prediction error is also consistent with the error observed on real genetic data when $p > n$, after adjusting for the effective number of independent markers. Unfortunately, due to computational issues, we were unable to perform the analysis in the $n \simeq p$ case with real data. However, we observed that the prediction accuracy already reaches almost 15% of the heritability of height when $n/p \simeq 5\%$, while De los Campos et al. (2013) suggested that its asymptotic upper bound is of the order of 20% of the heritability because of incomplete LD between causal loci and genotyped markers. Interestingly, ridge regression is not affected by correlated predictors, and consequently it is not affected by high LD between markers. When LD is high, this has the effect of reducing the degrees of freedom of the model (Dijkstra, 2014), which results in an improved prediction accuracy in comparison with a problem having the same number of independent predictors and the same heritability.

the empirical covariance matrices. Our approximation of the prediction accuracy in terms of correlation proposes a good compromise between existing approximations. In particular, it



Although our approximations and simulation results tend to show that the prediction accuracy can reach the heritability value when $n \gg p$, as already suggested by previous works (Daetwyler et al., 2008; de Vlaming and Groenen, 2015; Rabier et al., 2016), the large standard deviation of the prediction error that we observed between simulated individuals suggests that disease

risk prediction from genetic data alone is not accurate at the individual level, even for a relatively high heritability value in the context of a small additive effect hypothesis.

In direct continuity of this work, it would be interesting to investigate the behavior of prediction accuracy on real human data where $n \simeq p$. This would enable us to determine whether our approximations still hold in that case, and even in the case where $n > p$ (where we approximate the empirical covariance matrix of the markers to be diagonal). It would show whether it is possible for the prediction accuracy to exceed the upper bound proposed by De los Campos et al. (2013). A further prospect would be to consider a nonlinear model extension via kernel ridge regression, which may improve the prediction (Morota and Gianola, 2014).

DATA AVAILABILITY STATEMENT

The data analyzed in this study is subject to the following licenses/restrictions: There is a charge for access to the UK Biobank dataset. Requests to access these datasets should be directed to: <https://www.ukbiobank.ac.uk/>.

ETHICS STATEMENT

The studies involving human participants were reviewed and approved by NHS National Research Ethics Service North West (11/NW/0382). The patients/participants provided their written informed consent to participate in this study. Written informed consent was obtained from the individual(s) for the publication of any potentially identifiable images or data included in this article.

AUTHOR CONTRIBUTIONS

EL and CA designed and directed the study. AF, EL, and CA wrote the manuscript. AF created the synthetic datasets. AF performed all the analyses with substantial input from EL, CA, CD-R, and MP-J. All authors discussed the results and commented on the manuscript.

FUNDING

Financial support was obtained by the Laboratory of Excellence GENMED (Medical Genomics) grant no. ANR-10-LABX-0013 managed by the French National Research Agency (ANR) part of the Investment for the Future program.

ACKNOWLEDGMENTS

This research has been conducted using the UK Biobank Resource under application number 45408.

SUPPLEMENTARY MATERIAL

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fgene.2020.581594/full#supplementary-material>

REFERENCES

- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*. Information science and statistics. New York, NY: Springer.
- Brard, S., and Ricard, A. (2015). Is the use of formulae a reliable way to predict the accuracy of genomic selection? *J. Anim. Breed. Genet.* 132, 207–217. doi: 10.1111/jbg.12123
- Bulik-Sullivan, B. K., Loh, P.-R., Finucane, H. K., Ripke, S., Yang, J., Patterson, N., et al. (2015). LD score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nat. Genet.* 47, 291–295. doi: 10.1038/ng.3211
- Chen, G.-B. (2014). Estimating heritability of complex traits from genome-wide association studies using IBS-based Haseman-Elston regression. *Front. Genet.* 5:107. doi: 10.3389/fgene.2014.00107
- Daetwyler, H. D., Pong-Wong, R., Villanueva, B., and Woolliams, J. A. (2010). The impact of genetic architecture on genome-wide evaluation methods. *Genetics* 185, 1021–1031. doi: 10.1534/genetics.110.116855
- Daetwyler, H. D., Villanueva, B., and Woolliams, J. A. (2008). Accuracy of predicting the genetic risk of disease using a genome-wide approach. *PLoS ONE* 3:e3395. doi: 10.1371/journal.pone.0003395
- Dandine-Roulland, C., and Perdry, H. (2015). The use of the linear mixed model in human genetics. *Hum. Hered.* 80, 196–206. doi: 10.1159/000447634
- De los Campos, G., Vazquez, A. I., Fernando, R., Klimentidis, Y. C., and Sorensen, D. (2013). Prediction of complex human traits using the genomic best linear unbiased predictor. *PLoS Genet.* 9:e1003608. doi: 10.1371/journal.pgen.1003608
- de Vlaming, R., and Groenen, P. J. F. (2015). The current and future use of ridge regression for prediction in quantitative genetics. *BioMed Res. Int.* 2015:143712. doi: 10.1155/2015/143712
- Dijkstra, T. K. (2014). Ridge regression and its degrees of freedom. *Qual. Quant.* 48, 3185–3193. doi: 10.1007/s11135-013-9949-7
- Elsen, J.-M. (2017). An analytical framework to derive the expected precision of genomic selection. *Genet. Sel. Evol.* 49:95. doi: 10.1186/s12711-017-0366-6
- Feldman, M. W., and Lewontin, R. C. (1975). The heritability hang-up. *Science* 190, 1163–1168. doi: 10.1126/science.1198102
- Fisher, R. A. (1918). The correlation between relatives on the supposition of Mendelian inheritance. *Trans. R. Soc. Edinburgh* 52, 399–433. doi: 10.1017/S0080456800012163
- Ge, T., Chen, C.-Y., Neale, B. M., Sabuncu, M. R., and Smoller, J. W. (2017). Phenome-wide heritability analysis of the UK Biobank. *PLoS Genet.* 13:e1006711. doi: 10.1371/journal.pgen.1006711
- Goddard, M. (2009). Genomic selection: prediction of accuracy and maximisation of long term response. *Genetica* 136, 245–257. doi: 10.1007/s10709-008-9308-0
- Golan, D., Lander, E. S., and Rosset, S. (2014). Measuring missing heritability: inferring the contribution of common variants. *Proc. Natl. Acad. Sci. U.S.A.* 111, E5272–E5281. doi: 10.1073/pnas.1419064111
- Golub, G. H., Heath, M., and Wahba, G. (1978). Generalized cross-validation as a method for choosing a good ridge parameter. *Technometrics* 21, 215–233. doi: 10.1080/00401706.1979.10489751
- Henderson, C. R. (1975). Best linear unbiased estimation and prediction under a selection model. *Biometrics* 31, 423–447. doi: 10.2307/2529430
- Hirschhorn, J. N., and Daly, M. J. (2005). Genome-wide association studies for common diseases and complex traits. *Nat. Rev. Genet.* 6, 95–108. doi: 10.1038/nrg1521
- Hoerl, A. E., and Kennard, R. W. (1970). Ridge regression: biased estimation for nonorthogonal problems. *Technometrics* 12, 55–67. doi: 10.1080/00401706.1970.10488634
- Li, M.-X., Yeung, J. M., Cherny, S. S., and Sham, P. C. (2012). Evaluating the effective numbers of independent tests and significant p-value thresholds in commercial genotyping arrays and public imputation reference datasets. *Hum. Genet.* 131, 747–756. doi: 10.1007/s00439-011-1118-2
- Manolio, T. A., Collins, F. S., Cox, N. J., Goldstein, D. B., Hindorf, L. A., Hunter, D. J., et al. (2009). Finding the missing heritability of complex diseases. *Nature* 461, 747–753. doi: 10.1038/nature08494
- Meuwissen, T., Hayes, B., and Goddard, M. (2001). Prediction of total genetic value using genome-wide dense marker maps. *Genetics* 157, 1819–1829.
- Morota, G., and Gianola, D. (2014). Kernel-based whole-genome prediction of complex traits: a review. *Front. Genet.* 5:363. doi: 10.3389/fgene.2014.00363
- Patterson, H. D., and Thompson, R. (1971). Recovery of inter-block information when block sizes are unequal. *Biometrika* 58, 545–554. doi: 10.1093/biomet/58.3.545
- Perdry, H., and Dandine-Roulland, C. (2018). *gaston: Genetic Data Handling (QC, GRM, LD, PCA) & Linear Mixed Models. R package version 1.5.4*. Villejuif.
- Pharoah, P. D., Antoniou, A., Bobrow, M., Zimmern, R. L., Easton, D. F., and Ponder, B. A. (2002). Polygenic susceptibility to breast cancer and implications for prevention. *Nat. Genet.* 31, 33–36. doi: 10.1038/ng853
- Purcell, S., Wray, N., Stone, J., Visscher, P., O'Donovan, M. C., Sullivan, P. F., et al. (2009). Common polygenic variation contributes to risk of schizophrenia and bipolar disorder. *Nature* 460, 748–752. doi: 10.1038/nature08185
- Rabier, C.-E., Barre, P., Asp, T., Charmet, G., and Mangin, B. (2016). On the accuracy of genomic selection. *PLoS ONE* 11:e0156086. doi: 10.1371/journal.pone.0156086
- Robinson, G. K. (1991). That BLUP is a good thing: the estimation of random effects. *Stat. Sci.* 6, 15–32. doi: 10.1214/ss/1177011926
- Speed, D., and Balding, D. J. (2014). Multiblup: improved SNP-based prediction for complex traits. *Genome Res.* 24, 1550–1557. doi: 10.1101/gr.169375.113
- Visscher, P. M., and Goddard, M. E. (2015). A general unified framework to assess the sampling variance of heritability estimates using pedigree or marker-based relationships. *Genetics* 199, 223–232. doi: 10.1534/genetics.114.171017
- Wright, S. (1920). The relative importance of heredity and environment in determining the piebald pattern of guinea-pigs. *Proc. Natl. Acad. Sci. U. S. A.* 6:320. doi: 10.1073/pnas.6.6.320
- Wright, S. (1921). Correlation and causation. *J. Agric. Res.* 7, 557–585.
- Xu, S. (2003). Estimating polygenic effects using markers of the entire genome. *Genetics* 163, 789–801.
- Yang, J., Benyamin, B., McEvoy, B. P., Gordon, S., Henders, A. K., Nyholt, D. R., et al. (2010). Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* 42, 565–569. doi: 10.1038/ng.608
- Yang, J., Lee, S. H., Goddard, M. E., and Visscher, P. M. (2011a). GCTA: a tool for genome-wide complex trait analysis. *Am. J. Hum. Genet.* 88, 76–82. doi: 10.1016/j.ajhg.2010.11.011
- Yang, J., Manolio, T. A., Pasquale, L. R., Boerwinkle, E., Caporaso, N., Cunningham, J. M., et al. (2011b). Genome partitioning of genetic variation for complex traits using common SNPs. *Nat. Genet.* 43:519. doi: 10.1038/ng.823
- Zhao, B., and Zhu, H. (2019). Cross-trait prediction accuracy of high-dimensional ridge-type estimators in genome-wide association studies. *arXiv:1911.10142 [stat]*.
- Zhou, X. (2017). A unified framework for variance component estimation with summary statistics in genome-wide association studies. *Ann. Appl. Stat.* 11:2027. doi: 10.1214/17-AOAS1052

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Frouin, Dandine-Roulland, Pierre-Jean, Deleuze, Ambroise and Le Floch. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.