

Boosting Heritability: Estimating the Genetic Component of Phenotypic Variation With Multiple Sample Splitting

The Tien Mai (✉ t.t.mai@medisin.uio.no)

University of Oslo

Paul Turner

University of Oxford

Jukka Corander

University of Oslo

Research Article

Keywords: antimicrobial resistance, boosting, heritability, linear model.

Posted Date: February 17th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-200227/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published on March 27th, 2021. See the published version at <https://doi.org/10.1186/s12859-021-04079-7>.

Boosting heritability: estimating the genetic component of phenotypic variation with multiple sample splitting

The Tien Mai⁽¹⁾, Paul Turner^{(3),(4)} and Jukka Corander^{(1),(2)}

⁽¹⁾ Oslo Centre for Biostatistics and Epidemiology, Department of Biostatistics, University of Oslo, Norway.

⁽²⁾ Department of Mathematics and Statistics, University of Helsinki, Finland.

⁽³⁾ Cambodia-Oxford Medical Research Unit, Angkor Hospital for Children, Siem Reap, Cambodia.

⁽⁴⁾ Centre for Tropical Medicine and Global Health, Nuffield Department of Medicine, University of Oxford, UK.

Abstract

Background: Heritability is a central measure in genetics quantifying how much of the variability observed in a trait is attributable to genetic differences. Existing methods for estimating heritability are most often based on random-effect models, typically for computational reasons. The alternative of using a fixed-effect model has received much more limited attention in the literature.

Results: In this paper, we propose a generic strategy for heritability inference, termed as “*boosting heritability*”, by combining several advantageous features of different recent methods to produce an estimate of the heritability with a high-dimensional linear model. Boosting heritability uses in particular a multiple sample splitting strategy which leads to a more reliable estimate. We use both simulated data and real antibiotic resistance data from a major human pathogen, *Sptreptococcus pneumoniae*, to demonstrate the attractive features of our inference strategy.

Conclusions: Compared with other methods, boosting heritability yields a more reliable estimate and allows one to make inference about heritability.

Keywords: *antimicrobial resistance, boosting, heritability, linear model.*

*Corresponding author: t.t.mai@medisin.uio.no

1 Introduction

Whereas genome-wide association studies (GWAS) represent the primary tool for determining the genetic basis of any phenotype/trait of interest, quantifying the contribution of genetic factors to the variation of a phenotype plays in addition an important role in many studies. For this purpose, *heritability* is a crucial quantity [Falconer, 1960, Lynch and Walsh, 1998] and it is defined (in the narrow-sense) as the proportion of the variance of a phenotype explained by the (additive) genetic factors.

Current studies of heritability in the literature have usually been carried out in the linear mixed-effect model framework [Bulik-Sullivan et al., 2015, Yang et al., 2010]. In this framework, the effect sizes of genetic markers, usually SNPs, are assumed to be independent and identical distributed random variables, and often the normal distribution (with 0-mean) is used for computational reasons. The maximum likelihood and method of moments are the most widely used methods for heritability inference for this family of models [Yang et al., 2010, Golan et al., 2014, Bulik-Sullivan et al., 2015, Zhou, 2017, Bonnet, 2016, Speed et al., 2017].

Some comparisons of different methods for estimating heritability have been recently conducted, for example, in [Zhou, 2017, Evans et al., 2018, Weissbrod et al., 2018, Gorfine et al., 2017]. However, these works compare the performance of different methods on different datasets without paying much attention to the actual model specification. Since heritability is a concept detailing the additive variance of a trait which is in a certain sense based on a statistical model, heritability estimation is consequently dependent on the specified model [Zaitlen and Kraft, 2012]. For example, as reported in [Evans et al., 2018], there is a sizeable difference in the estimated heritability of schizophrenia \hat{h}_{SNP}^2 that equals 0.56 according to [Bulik-Sullivan et al., 2015] and only 0.23 according to [Lee et al., 2013]. These estimates have a very different interpretation also qualitatively and they disagree most likely because they are based on different statistical models of heritability.

In this paper, we focus on the high-dimensional linear regression model with fixed effects, where no distributional assumption on the effect sizes is made. Although limited from the computational perspective due to the extremely high-dimensional data in GWAS, high-dimensional linear regression is a natural model for GWAS in modelling the whole-genome level contributions of genetic variation. The benefit of this model over the classical univariate approach in GWAS has been demonstrated for example in [Wu et al., 2009, Brzyski et al., 2017]. The study of heritability estimation with fixed-effect models has been started relatively recently and it has not yet gained a wide-spread attention. A method of moments approach is proposed in [Dicker, 2014], a convex optimization strategy is suggested in [Janson et al., 2017] through a singular value decomposition, maximum likelihood estimation is studied in [Dicker and Erdogdu, 2016], and some adaptive procedures have also been theoretically studied in [Verzelen et al., 2018]. However, to our knowledge, a systematic numerical comparison of these

different methods for estimating heritability has not been made yet.

Some two-step procedures based on high-dimensional regularized regression have been introduced in [Gorfine et al., 2017, Li et al., 2019] that provide an insight to obtain more reliable and stable estimates of heritability. In brevity, this approach is based on splitting the data into two subsets. In the first step, variable selection is employed through a sparsity inducing regularization on one subset to select the relevant covariates. In the second step, these selected covariates are used to estimate heritability from the other subset of data. The selection step is useful to consider only a subset of the covariates that contribute to the variability of the trait (the response). Moreover, splitting the sample is done to avoid doing variable selection and heritability estimation on the same data which can cause overestimate [Li et al., 2019]. Although promising, this approach depends crucially on the particular partition used to split the data, which can lead to unstable estimates.

To achieve more reliable results, we propose to use a multiple sample splitting procedure so that different structures in the sample are presented in both selection and estimation steps with a sufficiently high probability [Meinshausen et al., 2009, Fan et al., 2012]. Based on this idea, we present a general framework called “boosting heritability” which allows a user to plug-in their own favourite method of variable selection and/or heritability estimation. By repeating sample splitting, one can also obtain various estimates of the heritability and thus provide a meaningful interval of the estimated values.

To demonstrate our framework, we apply the procedure to bacterial GWAS for estimating the heritability of antibiotic resistant phenotypes. While there are numerous works concerning estimating heritability in human GWAS, the topic has not yet been considered widely in bacteria, for the only prominent example see [Lees et al., 2017]. This is partly because bacterial GWAS poses unique challenges compare to studies with human or animal DNA, stemming from more limited recombination and highly structured populations that result in substantial linkage disequilibrium across whole chromosomes.

The paper is structured as follows. In Section 2 we present the linear model that relates a trait with a genotype matrix, then narrow-sense heritability is defined together with some discussion regarding the fixed-effect vs. random-effect approach for estimation. In Section 3, a set of methods for estimating heritability are defined, and subsequently we introduce our “boosting heritability” procedure. Results from a simulation study comparing the different methods as components of the framework presented in Section 4 and the application to antibiotic resistance phenotypes are presented in Section 5. We conclude and discuss our results in the final section.

2 Model and definition

Notations: Here, we introduce the main notations used in the paper. The ℓ_q norm ($0 < q < +\infty$) of a vector $x \in \mathbb{R}^d$ is defined by $\|x\|_q = (\sum_{i=1}^d |x_i|^q)^{1/q}$. For a matrix $A \in \mathbb{R}^{n \times m}$, $A_{i.}$ denotes its i -th row and $A_{.j}$ denotes its j -th column. For any index set $S \subseteq \{1, \dots, d\}$, x_S denotes the subvector of x containing only the components indexed by S , and A_S denotes the submatrix of A forming by columns of A indexed by S .

2.1 Model

Given a phenotype/trait y of n samples that is modelled as a linear combination of p genetic covariates $X_{.j}$ and an error term (environmental and unmeasured genetic effects)

$$y_i = X_{i.}\beta + \varepsilon_i, i = 1, \dots, n \quad (2.1)$$

where $X_{i.}$ are independent and identically distributed (i.i.d) with distribution $\mathcal{N}(0, \Sigma)$ and are independent of $\varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2)$.

Here we focus on the fixed effects encoded by β and assume that the genetic covariates X are random variables. Conversely, in the majority of works in the heritability literature assume that elements of β are considered as i.i.d. random variables following a Gaussian distribution i.e $\beta_j \stackrel{i.i.d}{\sim} \mathcal{N}(0, \sigma_\beta^2)$, while the genetic covariates X are assumed fixed.

2.2 Heritability

Under the model (2.1), we have for the i -th observation that

$$\text{Var}(y_i) = \text{Var}(X_{i.}\beta) + \sigma_\varepsilon^2 = \beta^\top \Sigma \beta + \sigma_\varepsilon^2.$$

We are interested in estimating (the narrow-sense) heritability for y defined as

$$h^2 = \frac{\beta^\top \Sigma \beta}{\beta^\top \Sigma \beta + \sigma_\varepsilon^2}. \quad (2.2)$$

Technically, heritability is a quantitative measure that expresses how much of the population variability present in a trait is due to genetic differences. Moreover, estimating heritability can assist in modelling the underlying genetic architecture. A heritability close to zero implies that environmental factors cause most of the variability of the trait. In contrast, a heritability close to 1 indicates that the variability of the trait is nearly exclusively caused by the differences in genetic factors.

As we have the relation

$$\mathbb{E}[\|y\|_2^2/n] = \text{Var}(y) = \beta^\top \Sigma \beta + \sigma_\varepsilon^2,$$

one can use $\|y\|_2^2/n$ as an unbiased estimator for the denominator of the heritability. Further, one can re-write (2.2) as

$$h^2 = 1 - \frac{\sigma_\varepsilon^2}{\text{Var}(y)} \tag{2.3}$$

and use an estimate of the noise-variance σ_ε^2 (see e.g. [Reid et al., 2016]) to estimate h^2 rather than directly estimate the genetic variance $\beta^\top \Sigma \beta$ (which requires an estimate of the covariance matrix and the effect sizes).

However, it is worth noting that as a bi-product from GWAS analysis when using a multivariate regression approach, such as the Elastic net discussed below, one would already have the estimated effect sizes corresponding to the selected covariates. Using these effect sizes to estimate the heritability would bring insight on the heritability corresponding to the selected covariates and thus clearly provide useful ways to understand the genetic architecture of a trait.

Contrasting the fixed and random effects

In GWAS the true number of causal loci reported tend to be comparatively small compared with the number of putative genetic markers p , which is usually in the order of hundreds of thousands at minimum. Assume that the true effect size β has $s \ll p$ non-zero entries. In the random-effect model, a further assumption is made concerning these non-zero entries such that they are i.i.d Gaussian $\mathcal{N}(0, \sigma_\beta^2)$. Under this random effect assumption, the heritability is defined [Bonnet et al., 2015, Li et al., 2019] as $s\sigma_\beta^2/(s\sigma_\beta^2 + \sigma_\varepsilon^2)$.

However, when employing the random-effect assumption, most methods do not use the sparsity constraint. This leads to the target heritability being estimating $p\sigma_\beta^2/(p\sigma_\beta^2 + \sigma_\varepsilon^2)$ and the resulting estimate of heritability may thus be inaccurate. Moreover, the LD structure, an important concept that represents the correlation structure of the covariates, is not directly addressed in the formula of heritability in random-effect model, which can make the estimate unjustifiable, e.g see [Speed et al., 2017, Speed and Balding, 2019]. Several attempts have been done recently to take into account the sparsity constraint within the random-effect model and some promising results have been reported in [Bonnet et al., 2015, Bonnet et al., 2018, Li et al., 2019].

3 Boosting heritability estimation

3.1 Related works and motivation

As the number of biomarkers can be very large, it is natural to first apply some variable selection or variable screening methods to remove the irrelevant variables from the actual heritability estimation phase. This kind of a post-selection approach has been proposed in the literature, more specifically for the fixed-effect model [Gorfine et al., 2017, Li et al., 2019].

The HERRA method proposed in [Gorfine et al., 2017] is based on a screening method (e.g. as in [Fan and Lv, 2008]) to reduce the number of covariates below the sample size. Given the remaining covariates, the sample is randomly divided into two equally sized parts. A lasso-type estimator is employed on the first subset to select a small number of important variables. After that, the least squares estimator is used on the second subset of data using only the selected covariates (from the lasso-type estimator) to get an estimate of the noise-variance. The role of the first and second subsets are switched to obtain another estimate of the noise-variance. Finally, heritability is calculated as in the formula (2.3) where the noise-variance is the mean of the two estimated noise-variances.

Another "two-stage" approach with sample-splitting has also been proposed in the paper [Li et al., 2019]. The data is randomly split into two disjoint equal sample size. On one half of the data, they use a sparse regularization method based on Elastic net to first select the relevant variables. Then, on the other half of the data, they only use the selected variables to estimate the heritability through a method of moments based approach [Dicker, 2014].

Both these approaches clearly suffer from some limitations. Firstly, when the number of covariates is very large, it is expensive to fit a sparse regularization directly as in the "two-stage" approach described above. Using a screening method, as in HERRA, to reduce the dimension of the problem is thus a pragmatic approach for applications. However, as the true number of causal biomarkers is not known, as well as their LD structure is not given, reducing the number of variables below the sample size (as in HERRA) introduces another problem from the practical perspective. Secondly, it is clear that both of these approaches crucially depend on the particular sample splitting employed. One can avoid this dependence by performing the sample splitting and inference procedure many times (e.g. 100 times) and aggregating the corresponding results [Meinshausen et al., 2009, Fan et al., 2012]. This is to ensure that the different latent structures possibly residing in the sample are properly taken into account in both the selection and estimation steps.

The idea of aggregating different estimates to yield an estimate with improved statistical properties is the central feature of the generic boosting approach widely used in

machine learning, such as AdaBoost [Freund and Schapire, 1996]. The multiple sample splitting approach has previously been proposed in statistics community as in [Meinshausen et al., 2009, Fan et al., 2012], and successfully used in GWAS [Renaux et al., 2020, Buzdugan et al., 2016].

3.2 Boosting heritability: multi sample splitting and aggregation of heritability

We propose a strategy that uses multiple sample splitting to estimate heritability, called Boosting heritability detailed in Algorithm 1.

Algorithm 1 Boosting heritability

- 1: **Step 0:** Using a screening method, such as a marginal-type sure independent screening [Fan and Lv, 2008], to remove 25% irrelevant covariates. This step aims at reduce the ultra-high dimension to a more manageable level.
 - 2: **Repeat** B times from step 1 to step 4,
 - 3: **Step 1:** With the remaining covariates, divide the sample uniformly at random into two equal parts.
 - 4: **Step 2:** On the first part of the data, use Elastic net to select the important covariates.
 - 5: **Step 3:** Then, on the second subset with only selected covariates from Step 2, estimate the heritability by using a method presented in Section 3.2.1.
 - 6: **Step 4:** Repeat Step 2 and Step 3 by changing the role of the first and second subset.
 - 7: **Final** \rightarrow The final heritability estimate is the mean of the estimated heritabilities at each repeat.
-

It is noted that the initial step (**Step 0**) is a screening step that can use a simple measure of association, such as the sample correlation, to remove covariates that are only weakly correlated with the trait of interest. This step is similar to the one used in HERRA [Gorfine et al., 2017] and in [Bonnet et al., 2018], however, we do not propose to reduce the number of covariates below the actual sample size. This is motivated by the fact for real data we do not know the true number of causal variates as well as the correlation structure of the variables. If too many covariates are removed, this can have a detrimental effect on the subsequent steps in the estimation procedure. Moreover, the initial screening step can be seen as optional, and necessary only for situations where the high dimensionality of the covariate space makes regularized model fitting tedious or practically impossible for practical purposes.

The sample splitting performed in **Step 1** is a useful method that can help to avoid overfitting when variable selection and subsequent estimation is considered [Fan et al., 2012, Buzdugan et al., 2016, Li et al., 2019]. **Step 2** corresponds to a variable

selection phase where we suggest to use Elastic Net as a default alternative, given its ability to deal with highly correlated covariates. Switching the roles of the data subsets help us to obtain a more stable estimate of the heritability. Note that by repeating sample splitting, various estimates of the heritability are obtained and thus provide a meaningful interval of the estimated values (for example see Figure 1).

From a computational point of view: The main cost for Boosting Heritability procedure is fitting a penalized regression (**Step 2**) for variable selection in the setting where $p \gg n$. However, fast computation methods for penalized regression on large GWAS data have been recently proposed see e.g [Qian et al., 2020]. Moreover, the B repetitions can easily be implemented in parallel.

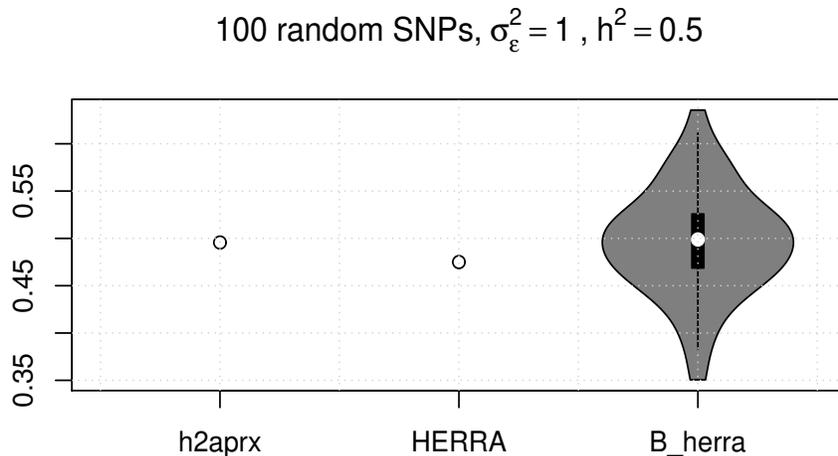


Figure 1: A violin plot for estimates of heritability from the simulation with MA data with 100 random covariates chosen as causal. We obtain an interval of heritability through the multiple sample splitting method (B_herra). See Section 4.

3.2.1 Plug-in Lasso type estimators for heritability

From the formula of heritability (2.2), direct approaches to estimate heritability can be obtained using estimates of the effect sizes β and of the covariance matrix. By using a lasso type method, one can obtain the non-zero estimated effect sizes of the selected covariates, and one can also use these covariates to obtain a sample covariance matrix. More precisely, let $S = \{j : \hat{\beta}_j \neq 0\}$ where $\hat{\beta}$ is an estimate from a lasso-type method, we can calculate the heritability as in equation (2.2) with $\Sigma_S = X_S X_S^\top / (n - 1)$,

$$\hat{h}^2 = \frac{\beta_S^\top \Sigma_S \beta_S}{\beta_S^\top \Sigma_S \beta_S + \sigma_\epsilon^2}.$$

The elastic net has been shown to be especially useful when the variables are dependent [Zou and Hastie, 2005] (LD structure), which is often the case with genetic marker data and this feature is especially highlighted in bacterial genome data. The corresponding estimator is defined as

$$\hat{\beta}_{Enet} := \arg \min_{\beta} \frac{1}{n} \sum_{i=1}^n \ell(y_i, \beta^T x_i) + \lambda [0.5(1 - \alpha) \|\beta\|_2^2 + \alpha \|\beta\|_1].$$

Here $\ell(a, b)$ is the negative log-likelihood for an observation e.g. for the linear Gaussian case it is $\frac{1}{2}(a - b)^2$ and for logistic regression it is $-a \cdot b + \log(1 + e^b)$. Elastic net is controlled by $\alpha \in [0, 1]$, that bridges the gap between lasso ($\alpha = 1$) and ridge regression ($\alpha = 0$). As the true genetic basis of a given trait is generally unknown as well as the LD structure is hard to estimate, we suggest to use a small value for α , e.g. 0.001. The tuning parameter $\lambda > 0$ controls the overall strength of the penalty and we use 10-fold cross-validation to choose suitable value for λ . Elastic net approach is implemented in the software 'pyseer' [Lees et al., 2018, Lees et al., 2020] focusing on GWAS for bacterial data.

4 Simulation studies

We use a real data set of 616 *Streptococcus pneumoniae* genomes collected from Massachusetts, denoted MA data, to create semi-synthetic datasets that incorporate levels of population structure and LD occurring in natural populations (see Figure 2). The data are publicly available through the publication [Croucher et al., 2015]. After initial data filtering with standard population genomic procedures (using a minor allele frequency threshold and removing missing data), we obtain a genotype matrix of 603 samples with 89703 SNPs. Using this observed genotype matrix, we simulate the responses/phenotypes through the linear model defined in (2.1).

Availability of data and code: The R codes and data used in the numerical experiments are available at: <https://github.com/tienmt/boostingher>.

4.1 Experimental designs

We consider the following designs for choosing the causal SNPs (non-zero effect sizes):

- sparse setting: 100 SNPs are randomly chosen.
- polygenic setting: 5000 SNPs are randomly chosen.
- Penicillin resistance ([Lees et al., 2016]): 100 SNPs are randomly chosen from 3 genes (pbpX, pbp1A, penA).

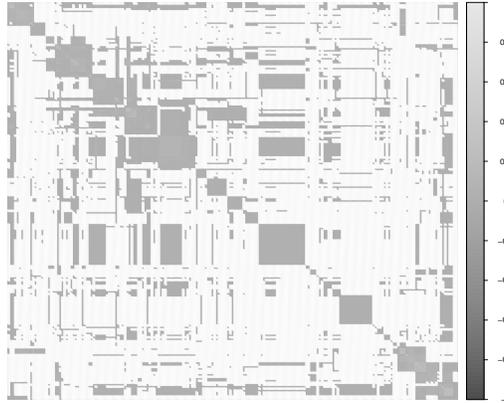


Figure 2: *Sample covariance matrix of the first 100 SNPs covariates in the genotype matrix shows the complex dependence structure present in the *S. pneumoniae* data.*

Then, the regression coefficients β^0 are sampled from the normal distribution $\mathcal{N}(0, 1)$. As the true covariance of the genotype matrix is not given, we need to re-normalize the coefficient β^0 as $\beta = \beta^0 \sqrt{\sigma_\varepsilon^2 h^2 / (\beta^{0\top} \bar{\Sigma} \beta^0 (1 - h^2))}$ to assure that the true corresponding heritability is approximating our target. Here h^2 is the target heritability and $\bar{\Sigma}$ is the sample covariance matrix of the genotype matrix. The target heritability is fixed as $h^2 = 0.5$ and we further examine the noise variance for each setup by altering between $\sigma_\varepsilon^2 = 1$ and $\sigma_\varepsilon^2 = 10^2$.

Note that we use the “oracle” estimator, denoted by ‘**h2aprx**’, that is calculated through the formula (2.3)

$$\text{h2aprx} = 1 - \frac{\sigma_\varepsilon^2}{\text{Var}(y)},$$

as a benchmark for comparison. As in simulations the true covariance matrix is not known in our setup, whereas the noise variance is given and thus this estimator provides a solid basis for approximating the true heritability.

For each setup, we generate 30 simulation runs and report the mean and the standard deviation of heritability estimates for each method across the simulation runs. We compare Elastic net (Enet), HERRA and the boosting versions of HERRA denoted by ‘B_herra’. The number of repeated sample splitting is performed with $B = 50$ times.

4.2 Results for estimating heritability

From the results in Table 1, it is clear that the “oracle” approximates well the target heritability in all designs. Generally, the boosting procedure tends to reduce the variability of the original underlying method it is used in conjunction with.

Elastic net underestimates the target, which can be explained by the downward bias known to influence the naive plug-in lasso-type approaches, such as the Elastic net. The effect is due to shrinkage of some of the coefficients corresponding to weak effect towards zero, while such weak effects may still be significant in terms of the total genetic trait variability. However, we would like to note that estimating heritability through Elastic net provides a good lower bound for the heritability, as indicated by the results.

On the other hand, HERRA and its boosting version return stable estimates. This can be anticipated as this approach follows the spirit of the ‘oracle’ estimator. More specifically, it aims at providing a consistent estimate of the noise variance and thus the corresponding heritability estimate would be also consistent and stable [Gorfine et al., 2017]. For this reason, the boosting HERRA will be our main focus method in real application in the next section.

Table 1: *Simulation results with MA data (standard deviation is given in parentheses).*

	100 causal SNPs, $\sigma_\varepsilon^2 = 1$	5000 causal SNPs, $\sigma_\varepsilon^2 = 1$	100 causal SNPs, $\sigma_\varepsilon^2 = 10^2$	100 causal SNPs from 3 genes, $\sigma_\varepsilon^2 = 1$
h2aprx	0.5004 (.0245)	0.5085 (.0256)	0.5026 (.0226)	0.4966 (.0227)
Enet	0.3585 (.0348)	0.3770 (.0500)	0.3724 (.0455)	0.3546 (.0386)
HERRA	0.5619 (.0507)	0.5583 (.0366)	0.5527 (.0494)	0.5204 (.0483)
B_herra	0.5551 (.0350)	0.5588 (.0294)	0.5467 (.0378)	0.5184 (.0371)

The effect of multiple data splitting

Clearly, choosing the number of data splittings B is crucial in practice. Here, we show that as B increases, the resulting estimated heritabilities concentrate around their mean, see Figure 3. Thus, we suggest to use at least $B \geq 30$ in practice and $B = 100$ would be a reasonable choice, computational resources permitting.

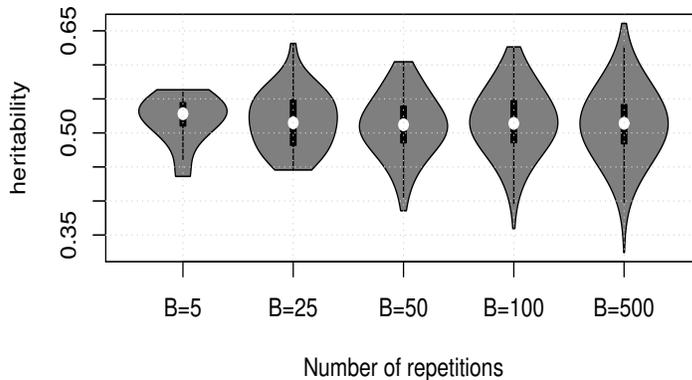


Figure 3: *Simulation results with MA data, 100 randomly selected SNPs, $\sigma_\epsilon^2 = 1$ and the target heritability $h^2 = 0.5$. Violin plot depicts the distribution of heritability estimates for choosing B , the number of repeating data splittings.*

The effect of the screening step

We further investigate the effect of reducing the covariates by using the screening step. Different scenarios for 100 randomly selected SNPs with target heritability $h^2 = 0.5$ and $\sigma_\epsilon^2 = 1$ are examined, see Table 2. More precisely, we consider two scenarios: remove 90% of the covariates, and only retain top $n + 1$ covariates.

It is revealed that using the screening step to reduce the irrelevant covariates not only reduces the dimension of the data, but can also improve the heritability estimation, for the scenario of removing 90% of the covariates. This fact has also been reported before in the linear mixed model approach in [Bonnet et al., 2018], where the authors show an improvement of the maximum likelihood estimation. However, if too many covariates are removed, heritability estimation can be inaccurate as in the scenario of keeping only top $n + 1$ covariates.

Table 2: *Simulation results with MA data, 100 randomly selected SNPs, $\sigma_\epsilon^2 = 1$ and $h^2 = 0.5$ (standard deviation is given in parentheses).*

	remove 90% covariates	keep $(n + 1)$ covariates
h2aprx	0.5010 (.0226)	0.4988 (.0229)
Enet	0.4428 (.0337)	0.3262 (.0189)
HERRA	0.4788 (.0384)	0.4063 (.0270)
B_herra	0.4740 (.0318)	0.4046 (.0274)

5 Heritability of antibiotic resistance in Maela data

To further illustrate the boosting based approach, we apply our procedure to Maela data which represent 3069 *Streptococcus pneumoniae* genomes from an infant cohort study conducted in a refugee camp on the Thailand-Myanmar border [Chewapreecha et al., 2014, Lees et al., 2016]. After some data filtering with standard population genomic procedures (using a minor allele frequency threshold and removing missing data), we obtain a genotype matrix with 121014 SNPs. We consider resistances to five different antibiotics as the phenotypes: chloramphenicol, erythromycin, tetracycline, penicillin and co-trimoxazole.

The heritability of the antibiotic resistance phenotype is expected to be high, meaning that the variability stems primarily from the observed genetic differences among these bacteria. We use two different types of resistance phenotypes to investigate their heritability. First we use the binary phenotype corresponding to the labels 'R' or 'S' (stand for 'Sensitive' or 'Resistant') for each bacterial isolate in the cohort. Second, we use a continuous phenotype corresponding to the inhibition zone diameters measured in the lab. These inhibition zone diameters are in practice used to defined whether a sample is sensitive or resistant to an antibiotic. It is however worthwhile noting that the transformation from inhibition zone diameters to labelling a sample 'S' vs 'R' is nonlinear due to the way the inhibition mechanism dynamics in the bacterial culture.

We apply Enet, HERRA, boosting version of HERRA and GCTA method to this data. More specifically, GCTA [Yang et al., 2011] is a widely used method based on a linear mixed model and maximum (restricted) likelihood estimation. The results are given in the Table 3 and Table 4 for the two data types, respectively.

Table 3: *Heritabilities of antibiotic resistance (binary) phenotypes in Maela data (standard deviation is given in parentheses).*

	Enet	HERRA	B_herra	GCTA
Chloramphenicol	0.4623	0.7489	0.7617 (.0413)	0.8257 (.0132)
Erythromycin	0.7979	0.9150	0.9140 (.0119)	0.7990 (.0141)
Tetracycline	0.8217	0.8899	0.8928 (.0113)	0.8260 (.0127)
Penicillin	0.7369	0.8237	0.8280 (.0138)	0.6695 (.0228)
Co-trimoxazole	0.5324	0.6093	0.6340 (.0368)	0.6005 (.0249)

Table 4: *Heritabilities of antibiotic resistance phenotypes using inhibition zone diameters in Maela data (standard deviation is given in parentheses).*

	Enet	HERRA	B_herra	GCTA
Chloramphenicol	0.5133	0.6364	0.6337 (.0267)	0.6837 (.0226)
Erythromycin	0.7350	0.8413	0.8383 (.0140)	0.7282 (.0196)
Tetracycline	0.7364	0.8072	0.8435 (.0135)	0.7514 (.0178)
Penicillin	0.8092	0.8445	0.8462 (.0132)	0.7123 (.0202)
Co-trimoxazole	0.7104	0.7840	0.7571 (.0210)	0.7826 (.0157)

As a broad summary, heritabilities of these five antibiotic resistances are high, as expected, whether using binary or continuous phenotypes. However, we would like to note that the results for binary responses are on the observed scale (0/1 resistance status), as we are not able to transform them into the underlying threshold model, see [Gorfine et al., 2017]. The Elastic net method yields an important insight by providing a lower bound on the heritability of these antibiotic resistances. For continuous phenotypes, it is at least 51% for chloramphenicol, at least 73% for erythromycin, at least 73% for tetracycline, at least 80% for penicillin and at least 71% for co-trimoxazole.

Interestingly, B_herra yields consistent results with GCTA method. However, the result for heritability of penicillin by GCTA is lower than the one from Enet method while boosting HERRA is not, see Table 4.

6 Discussion and conclusions

In this paper, we provide a general framework 'boosting heritability' for making inference about heritability. The main ingredient of 'boosting heritability' is a multiple sample splitting strategy. This strategy allows one to employ a variable selection step to remove irrelevant covariates that do not contribute to the variability of a trait and thus produce a reliable estimate of heritability. Moreover, by repeating sample splitting many times, this strategy makes sure that different latent structures are taken into account in both selection and estimation steps.

Numerical comparisons of different methods together with our proposal for estimating heritability in linear (fixed-effect) model draw a systematic picture on the behaviour of the current approaches when focusing on an application to bacterial GWAS. The results on real data suggest that the observed variability of the five studied antibiotic resistances is mainly due to the variability in the observed genetic factors, while some unexplained variation still remains.

Succeeding in improving and stabilizing HERRA [Gorfine et al., 2017], "boosting heritability" framework still preserves its advantages that are able to deal with the

dichotomous, time-to-event or age-at-onset traits. Moreover, boosting heritability procedure is also applicable for random-effect model where the heritability estimation step (Step 3 in Algorithm 1) is done by using a random effect method as in [Li et al., 2019]. These would be possible new research directions for the future.

Furthermore, our boosting heritability procedure uses a simple aggregation to combine the estimates that is to use their arithmetic mean. Other types of aggregation, see e.g [Renaux et al., 2020, Buzdugan et al., 2016], could also be used and further examined in future works.

Acknowledgments

T.T.M and J.C. would like to thank John A Lees for his useful discussion on GWAS and heritability.

Funding

This research was supported by the European Research Council grant no. 742158.

Availability of data and materials

The R codes and data used in the numerical experiments are available at: <https://github.com/tienmt/boostingher> .

Ethics approval and consent to participate

Not applicable.

Competing interests

The authors declare that they have no competing interests

Consent for publication

Not applicable.

Authors' contributions

Conceptualization: T.T.M. and J.C. . Formal analysis: T.T.M.. Data curation: P.T. . Methodology, T.T.M.. Writing: original draft, T.T.M.; review and editing, all authors. Funding acquisition: J.C..

References

- Bonnet, A. (2016). *Heritability Estimation in High-dimensional Mixed Models: Theory and Applications*. PhD thesis, Université Paris-Saclay.
- Bonnet, A., Gassiat, E., Lévy-Leduc, C., et al. (2015). Heritability estimation in high dimensional sparse linear mixed models. *Electronic Journal of Statistics*, 9(2):2099–2129.
- Bonnet, A., Lévy-Leduc, C., Gassiat, E., Toro, R., and Bourgeron, T. (2018). Improving heritability estimation by a variable selection approach in sparse high dimensional linear mixed models. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 67(4):813–839.
- Brzyski, D., Peterson, C. B., Sobczyk, P., Candès, E. J., Bogdan, M., and Sabatti, C. (2017). Controlling the rate of gwas false discoveries. *Genetics*, 205(1):61–75.
- Bulik-Sullivan, B. K., Loh, P.-R., Finucane, H. K., Ripke, S., Yang, J., Patterson, N., Daly, M. J., Price, A. L., Neale, B. M., of the Psychiatric Genomics Consortium, S. W. G., et al. (2015). Ld score regression distinguishes confounding from polygenicity in genome-wide association studies. *Nature genetics*, 47(3):291.
- Buzdugan, L., Kalisch, M., Navarro, A., Schunk, D., Fehr, E., and Bühlmann, P. (2016). Assessing statistical significance in multivariable genome wide association analysis. *Bioinformatics*, 32(13):1990–2000.
- Chewapreecha, C., Marttinen, P., Croucher, N. J., Salter, S. J., Harris, S. R., Mather, A. E., Hanage, W. P., Goldblatt, D., Nosten, F. H., Turner, C., et al. (2014). Comprehensive identification of single nucleotide polymorphisms associated with beta-lactam resistance within pneumococcal mosaic genes. *PLoS genetics*, 10(8):e1004547.

- Croucher, N. J., Finkelstein, J. A., Pelton, S. I., Parkhill, J., Bentley, S. D., Lipsitch, M., and Hanage, W. P. (2015). Population genomic datasets describing the post-vaccine evolutionary epidemiology of streptococcus pneumoniae. *Scientific data*, 2:150058.
- Dicker, L. H. (2014). Variance estimation in high-dimensional linear models. *Biometrika*, 101(2):269–284.
- Dicker, L. H. and Erdogdu, M. A. (2016). Maximum likelihood for variance estimation in high-dimensional linear models. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, PMLR 51:159-167, 2016*.
- Evans, L. M., Tahmasbi, R., Vrieze, S. I., Abecasis, G. R., Das, S., Gazal, S., Bjelland, D. W., Candia, T. R., Goddard, M. E., Neale, B. M., et al. (2018). Comparison of methods that use whole genome data to estimate the heritability and genetic architecture of complex traits. *Nature genetics*, 50(5):737.
- Falconer, D. S. (1960). *Introduction to quantitative genetics*. Oliver And Boyd; Edinburgh; London.
- Fan, J., Guo, S., and Hao, N. (2012). Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(1):37–65.
- Fan, J. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(5):849–911.
- Freund, Y. and Schapire, R. E. (1996). Experiments with a new boosting algorithm. In *Proceedings of the Thirteenth International Conference on Machine Learning, ICML’96*, pages 148–156.
- Golan, D., Lander, E. S., and Rosset, S. (2014). Measuring missing heritability: inferring the contribution of common variants. *Proceedings of the National Academy of Sciences*, 111(49):E5272–E5281.
- Gorfine, M., Berndt, S. I., Chang-Claude, J., Hoffmeister, M., Le Marchand, L., Potter, J., Slattery, M. L., Keret, N., Peters, U., and Hsu, L. (2017). Heritability estimation using a regularized regression approach (herra): Applicable to continuous, dichotomous or age-at-onset outcome. *PloS one*, 12(8):e0181269.
- Janson, L., Barber, R. F., and Candes, E. (2017). Eigenprism: inference for high dimensional signal-to-noise ratios. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1037–1065.
- Lee, S. H., Ripke, S., Neale, B. M., Faraone, S. V., Purcell, S. M., Perlis, R. H., Mowry, B. J., Thapar, A., Goddard, M. E., Witte, J. S., et al. (2013). Genetic relationship

- between five psychiatric disorders estimated from genome-wide snps. *Nature genetics*, 45(9):984.
- Lees, J. A., Croucher, N. J., Goldblatt, D., Nosten, F., Parkhill, J., Turner, C., Turner, P., and Bentley, S. D. (2017). Genome-wide identification of lineage and locus specific variation associated with pneumococcal carriage duration. *Elife*, 6:e26255.
- Lees, J. A., Galardini, M., Bentley, S. D., Weiser, J. N., and Corander, J. (2018). pyseer: a comprehensive tool for microbial pangenome-wide association studies. *Bioinformatics*, 34(24):4310–4312.
- Lees, J. A., Mai, T. T., Galardini, M., Wheeler, N. E., Horsfield, S. T., Parkhill, J., and Corander, J. (2020). Improved prediction of bacterial genotype-phenotype associations using interpretable pangenome-spanning regressions. *Mbio*, 11(4).
- Lees, J. A., Vehkala, M., Välimäki, N., Harris, S. R., Chewapreecha, C., Croucher, N. J., Marttinen, P., Davies, M. R., Steer, A. C., Tong, S. Y., et al. (2016). Sequence element enrichment analysis to determine the genetic basis of bacterial phenotypes. *Nature communications*, 7:12797.
- Li, X., Wu, D., Cui, Y., Liu, B., Walter, H., Schumann, G., Li, C., and Jiang, T. (2019). Reliable heritability estimation using sparse regularization in ultrahigh dimensional genome-wide association studies. *BMC bioinformatics*, 20(1):219.
- Lynch, M. and Walsh, B. (1998). *Genetics and analysis of quantitative traits*, volume 1. Sinauer Sunderland, MA.
- Meinshausen, N., Meier, L., and Bühlmann, P. (2009). P-values for high-dimensional regression. *Journal of the American Statistical Association*, 104(488):1671–1681.
- Qian, J., Tanigawa, Y., Du, W., Aguirre, M., Chang, C., Tibshirani, R., Rivas, M. A., and Hastie, T. (2020). A fast and scalable framework for large-scale and ultrahigh-dimensional sparse regression with application to the uk biobank. *PLoS genetics*, 16(10):e1009141.
- Reid, S., Tibshirani, R., and Friedman, J. (2016). A study of error variance estimation in lasso regression. *Statistica Sinica*, 26:35–67.
- Renaux, C., Buzdugan, L., Kalisch, M., and Bühlmann, P. (2020). Hierarchical inference for genome-wide association studies: a view on methodology with software. *Computational Statistics*, 35(1):1–40.
- Speed, D. and Balding, D. J. (2019). Sumher better estimates the snp heritability of complex traits from summary statistics. *Nature genetics*, 51(2):277.
- Speed, D., Cai, N., Johnson, M. R., Nejentsev, S., Balding, D. J., Consortium, U., et al. (2017). Reevaluation of snp heritability in complex human traits. *Nature genetics*, 49(7):986.

- Verzelen, N., Gassiat, E., et al. (2018). Adaptive estimation of high-dimensional signal-to-noise ratios. *Bernoulli*, 24(4B):3683–3710.
- Weissbrod, O., Flint, J., and Rosset, S. (2018). Estimating snp-based heritability and genetic correlation in case-control studies directly and with summary statistics. *The American Journal of Human Genetics*, 103(1):89–99.
- Wu, T. T., Chen, Y. F., Hastie, T., Sobel, E., and Lange, K. (2009). Genome-wide association analysis by lasso penalized logistic regression. *Bioinformatics*, 25(6):714–721.
- Yang, J., Benyamin, B., McEvoy, B. P., Gordon, S., Henders, A. K., Nyholt, D. R., Madden, P. A., Heath, A. C., Martin, N. G., Montgomery, G. W., et al. (2010). Common snps explain a large proportion of the heritability for human height. *Nature genetics*, 42(7):565.
- Yang, J., Lee, S. H., Goddard, M. E., and Visscher, P. M. (2011). Gcta: a tool for genome-wide complex trait analysis. *The American Journal of Human Genetics*, 88(1):76–82.
- Zaitlen, N. and Kraft, P. (2012). Heritability in the genome-wide association era. *Human genetics*, 131(10):1655–1664.
- Zhou, X. (2017). A unified framework for variance component estimation with summary statistics in genome-wide association studies. *The annals of applied statistics*, 11(4):2027.
- Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)*, 67(2):301–320.

Figures

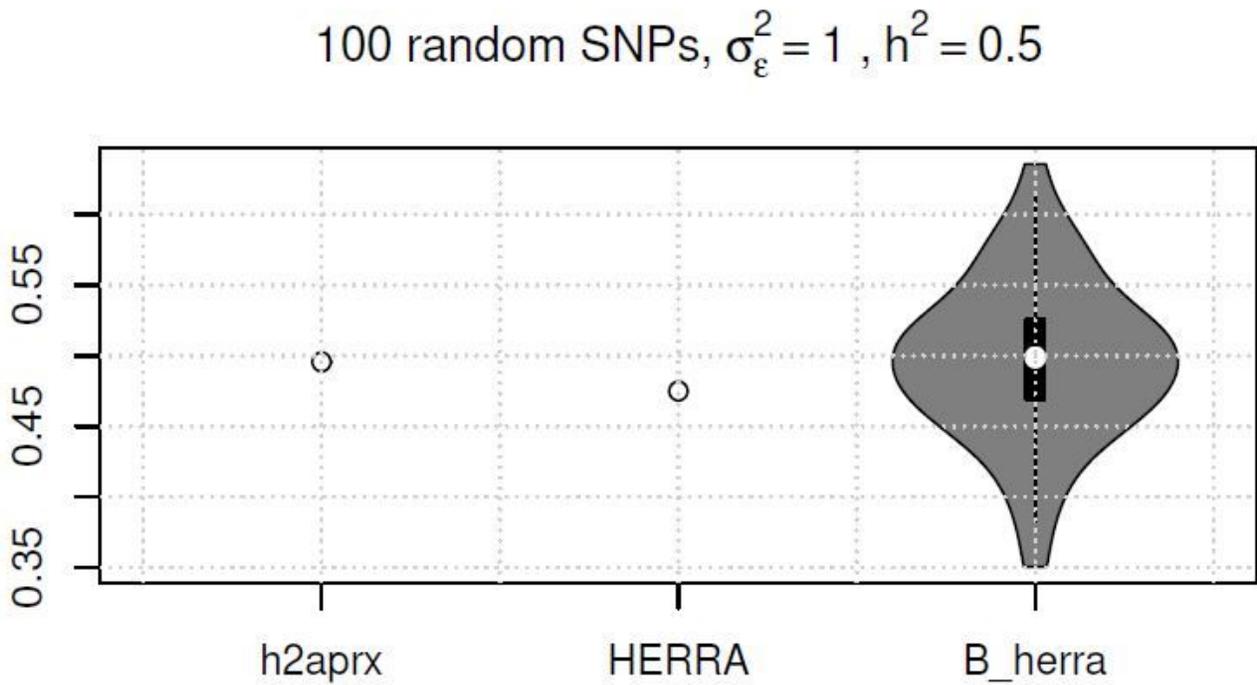


Figure 1

A violin plot for estimates of heritability from the simulation with MA data with 100 random covariates chosen as causal. We obtain an interval of heritability through the multiple sample splitting method (B_herra). See Section 4.

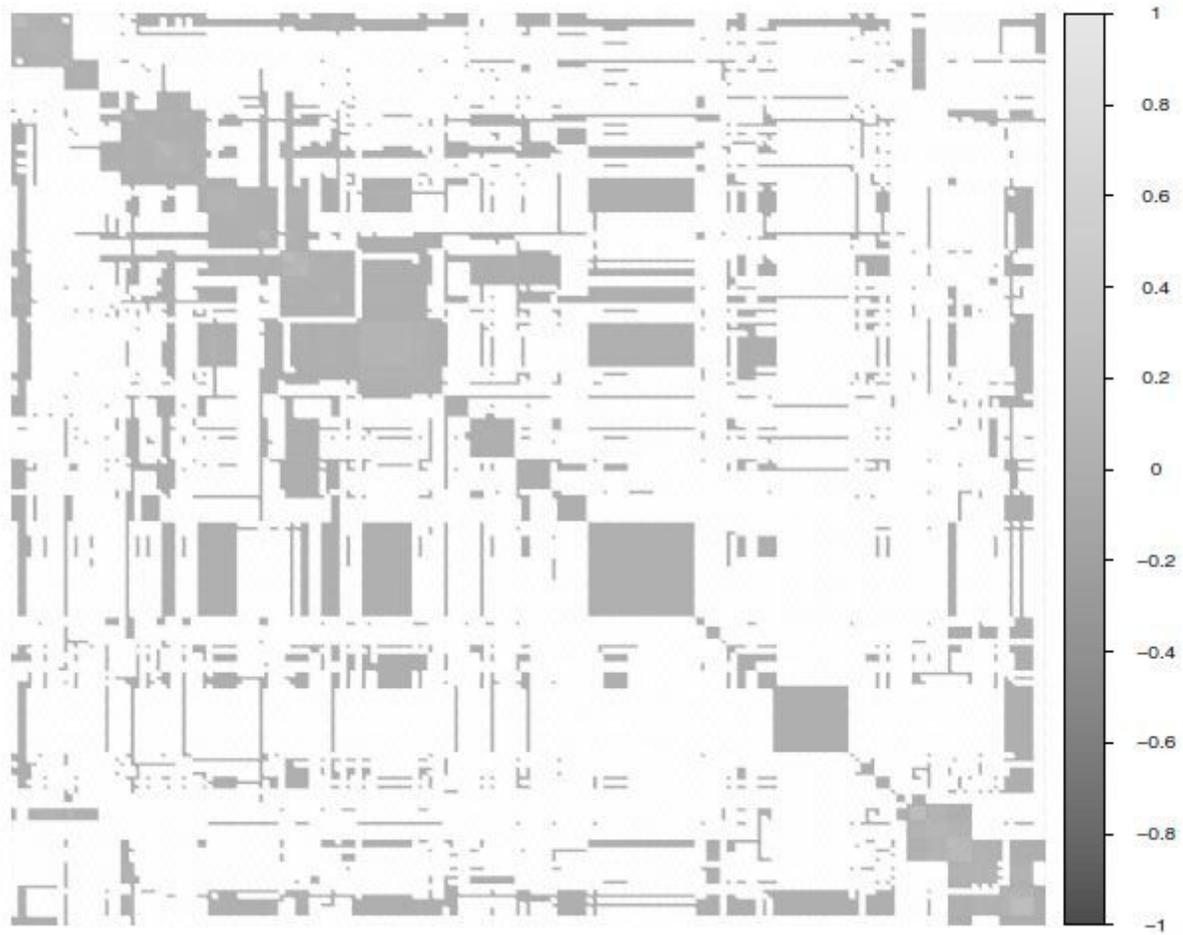


Figure 2

Sample covariance matrix of the first 100 SNPs covariates in the genotype matrix shows the complex dependence structure present in the *S. pneumoniae* data.

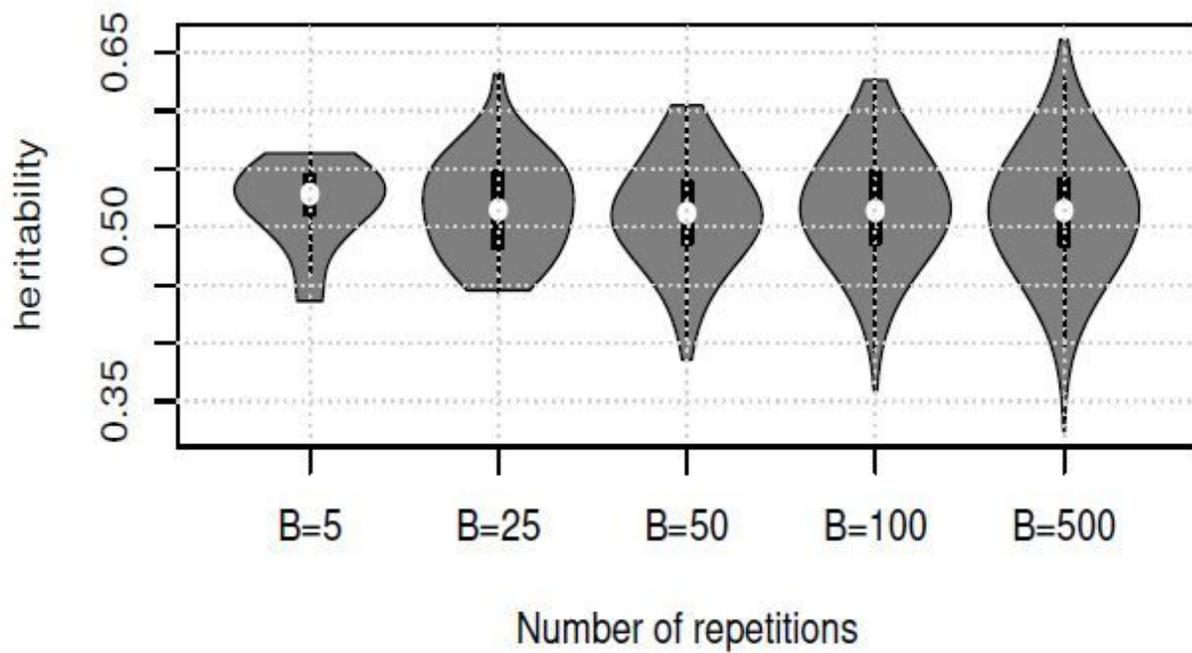


Figure 3

"Please see the Manuscript PDF file for the complete figure caption".