

An evaluation of the interpretability and predictive performance of the BayesR model for genomic prediction

Fanny Mollandin^{*,1}, Andrea Rau^{*,†,1} and Pascal Croiseau^{*,1}

^{*}Université Paris-Saclay, INRAE, AgroParisTech, GABI, [†]BioEcoAgro Joint Research Unit, INRAE, Université de Liège, Université de Lille, Université de Picardie Jules Verne

ABSTRACT Technological advances and decreasing costs have led to the rise of increasingly dense genotyping data, making feasible the identification of potential causal markers. Custom genotyping chips, which combine medium-density genotypes with a custom genotype panel, can capitalize on these candidates to potentially yield improved accuracy and interpretability in genomic prediction. A particularly promising model to this end is BayesR, which divides markers into four effect size classes. BayesR has been shown to yield accurate predictions and promise for quantitative trait loci (QTL) mapping in real data applications, but an extensive benchmarking in simulated data is currently lacking. Based on a set of real genotypes, we generated simulated data under a variety of genetic architectures, phenotype heritabilities, and we evaluated the impact of excluding or including causal markers among the genotypes. We define several statistical criteria for QTL mapping, including several based on sliding windows to account for linkage disequilibrium. We compare and contrast these statistics and their ability to accurately prioritize known causal markers. Overall, we confirm the strong predictive performance for BayesR in moderately to highly heritable traits, particularly for 50k custom data. In cases of low heritability or weak linkage disequilibrium with the causal marker in 50k genotypes, QTL mapping is a challenge, regardless of the criterion used. BayesR is a promising approach to simultaneously obtain accurate predictions and interpretable classifications of SNPs into effect size classes. We illustrated the performance of BayesR in a variety of simulation scenarios, and compared the advantages and limitations of each.

KEYWORDS
genomic prediction
QTL mapping
Bayesian model

INTRODUCTION

The primary objective of genomic prediction is to use genomic variation, usually single nucleotide polymorphisms (SNPs), to predict phenotypes, i.e. an observable trait of an individual. In particular, genomic prediction models are widely used as an evaluation tool

for genomic selection in animal breeding (1), and for the calculation of polygenic risk scores for human diseases (2). As genotyping costs have declined (3), there has been a corresponding increase in the amount of genotyping data available for analysis. In addition, lower costs and better data storage capacity have allowed for increasingly dense genotypes, up to and including whole genome sequences (WGS), which in turn have enabled sequence-level genotypes to be imputed for individuals genotyped using lower density chips(4). However, analyzing these increasingly large genotype data can come at a high computational cost and requires suitable statistical methods. Although the use of higher density genotypes was initially thought to hold promise for improved prediction accuracy, their performance was not found to improve that of high density chips in real data, due to the inclusion of a large number of

Manuscript compiled: Wednesday 21st October, 2020

[†]These authors contributed equally to this work.

¹Fanny Mollandin

Université Paris-Saclay, INRAE, AgroParisTech, GABI

Allée de Vilvert

78350, Jouy-en-Josas

fanny.mollandin@inrae.fr

1 non-causative SNPs (5). While the exhaustive use of WGS variants
2 has not led to meaningful improvements in prediction, they do al-
3 low for the direct inclusion of candidate, or even causal, mutations
4 (6). For simplicity, we refer to such mutations as quantitative trait
5 loci (QTL) throughout. If such QTLs are known a priori or can be
6 directly identified through variable selection in the model itself,
7 this could potentially lead to the double advantage of improv-
8 ing both the accuracy and interpretability of genomic prediction
9 models(7; 8). With this in mind, custom chips, which include SNPs
10 from a medium-density chip (intended to cover the genome) as
11 well as candidates or causal mutations for a set of traits, have
12 been developed, offering the cost and computational advantages
13 of a reasonably sized chip with the increased predictive ability
14 and interpretability provided by the inclusion of potential causal
15 mutations.

16 Most models used in routine genomic selection are based on
17 linear models, notably best linear unbiased prediction (BLUP)
18 and genomic BLUP (GBLUP). These models assume that all SNPs
19 contribute equally to the genomic variance, with each SNP effect
20 following a normal distribution with common variance. Although
21 the assumption about common SNP effects allows for great com-
22 putational efficiency, it is quite strong and can limit the biological
23 interpretability of results. To address this limitation, although
24 deep learning models have recently started to appear (9; 10), a
25 more frequent alternative is the set of non-linear Bayesian models
26 comprising the so-called Bayesian alphabet. These include, among
27 others, BayesA (1), BayesB (1), BayesC π (11), BayesR (12), and
28 BayesRC (13). The aim of all of these models is to improve pre-
29 dictive accuracy by better estimating SNP effects through more
30 flexible prior specifications. For instance, in the earliest model
31 introduced, BayesA, all markers are assumed to be drawn from a
32 normal distribution whose variance follows an $\text{Inv}-\chi^2$ distribu-
33 tion. Although the assumptions of BayesA are arguably closer to
34 reality than BLUP or GBLUP, it is computationally expensive to es-
35 timate variances for every SNP in dense genotyping data. Instead,
36 a useful alternative is to assume that a (potentially large) portion
37 of markers contribute no genetic variance. This is the strategy
38 employed by both BayesB and BayesC, which model marker effect
39 variances as a zero-inflated distribution by assigning null effects
40 with a fixed probability, and assuming the variance of non-null
41 SNPs respectively follow a per-SNP or common $\text{Inv}-\chi^2$ distribu-
42 tion. BayesC π further assumes that the proportion of null SNP
43 effects is itself a random variable, and otherwise uses a common
44 prior distribution for non-null SNP effects. BayesR provides addi-
45 tional flexibility by defining four classes of SNP effect size (null,
46 small, medium, large), where SNP effects are modeled using a four-
47 component normal mixture model. The related BayesRC model
48 further allows for SNPs to be grouped into disjoint categories (e.g.,
49 according to prior biological information), for which the BayesR
50 model is subsequently fit independently.

51 Although these Bayesian genomic prediction models are mainly
52 used for phenotype prediction, they also provide valuable per-
53 SNP information, including posterior estimates of effect size and
54 variance, which could be used for QTL mapping. In contrast to
55 genome-wide association study (GWAS) methods, SNP effects
56 are estimated simultaneously and make use of variable selection
57 within the model itself, rather than relying on univariate hypothe-
58 sis tests and corrections for multiple testing. As the quantity and
59 quality of prior biological knowledge continues to improve and
60 the identification of causal mutations from WGS data (14) becomes
61 increasingly feasible, the flexible model definition of BayesR and
62 BayesRC thus make them interesting candidates for simultane-

ously providing good predictability and biologically interpretable
QTL mapping results. In this spirit, Moser *et al.* showed encourag-
ing results for the use of BayesR in complex traits for prediction and
QTL mapping in real data (15). However, a comprehensive sim-
ulation study investigating the interpretability and performance
of BayesR in a wide variety of settings is currently lacking in the
literature. In addition, to date there has been little discussion of
the various criteria that can potentially be used to map QTLs using
the BayesR model output.

To address this gap, our goal in this work is to identify the
coherence between the BayesR model specification and known
QTL effects in simulated data under a variety of conditions. The
BayesR approach is of particular interest here, as it has been shown
in the literature to improve prediction accuracy (16), but its ability
to correctly assign QTLs to the appropriate effect size categories
has not yet been extensively evaluated in simulations. We focus on
the case where a prior categorization of markers (i.e., the BayesRC
approach) is not available. Using simulated data, we evaluate the
robustness of BayesR under a wide variety of genetic architectures,
phenotype heritabilities, and polygenic variances, and we illustrate
the conditions under which BayesR successfully identifies known
QTLs while maintaining high accuracy for phenotypic prediction.
Finally, we describe and compare several statistical criteria that
can be used to perform QTL mapping using BayesR output. Based
on the results of our simulation study, we discuss the optimal
framework for an accurate and interpretable analysis using BayesR,
as well as its limitations.

MATERIALS AND METHODS

Data simulation based on real genotypes

To maintain a realistic linkage disequilibrium (LD) structure among
SNPs, we generated simulated data based on a set of genotypes
assayed using Illumina Bovine SNP50 BeadChip arrays from $n =$
2605 Montbéliarde bulls. We divided individuals into learning
and validation sets (i.e., the “holdout method”), with the 80%
oldest bulls ($n_{\text{learning}} = 2083$) in the former and the 20% youngest
($n_{\text{validation}} = 522$) in the latter to reflect the strategy typically used
in routine genomic selection. We excluded SNPs with a minor
allele frequency (MAF) less than 0.01, leaving a total of $p = 46,178$
SNPs.

To simulate phenotypes \mathbf{y} for the $n = 2605$ bulls, we made use
of a standard linear model:

$$\mathbf{y} = \mu \mathbf{1}_n + \mathbf{X}\beta + e, \quad (1)$$
$$e \sim N(0, \mathbf{I}_n \sigma_e^2)$$

where μ denotes the trait mean (including fixed effects), β the vec-
tor of effects for the p SNPs, \mathbf{X} the centered and scaled genotype
design matrix, and e the residuals, assumed to follow a normal
distribution with variance σ_e^2 . Parameters for this linear model
were set as follows. For each simulated dataset we sampled from
the available SNPs a set of n_{QTL} QTLs and a set of n_{poly} polygenic
SNPs, as well as their corresponding genetic variances for each
selected marker. To reduce the impact of extreme MAFs on ge-
nomic prediction (17) and QTL detection, we focused on frequent
QTLs by drawing the n_{QTL} and n_{poly} SNPs from those with a MAF
 ≥ 0.15 . In all simulations, we selected a total of $n_{\text{QTL}} = 5$ large
QTLs, varying the corresponding proportion k of total genetic ad-
ditive variance σ_g^2 as described below. The phenotypic variance
and mean were respectively set to $\sigma_y^2 = 100$ and $\mu = 0$, and SNP
heritability $h^2 = \frac{\sigma_g^2}{\sigma_y^2}$ was varied across simulation settings.

Number of QTLs	5	5	5	5	5	5	5	5	5	5	5	5	5
Number of polygenic SNPs	9637	9550	9500	9450	9350	9250	9100	9000	8750	8500	8250	8000	7500
Per-QTL % of σ_g^2	0.725	0.9	0.10	0.11	0.13	0.15	0.18	2	2.5	3	3.5	4	5
Per-polygenic SNP % of σ_g^2	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01

Table 1 Simulation settings for each of the 13 QTL effect-size scenarios considered for each given level of heritability, $h^2 = \{0.1, 0.3, 0.5, 0.8\}$. The number of simulated QTLs, number of polygenic SNPs, percentage of genetic variance attributed to each QTL, and percentage of genetic variance attributed to each polygenic SNP are provided. Summing the percentage of genetic variance explained by the total number of QTLs and polygenic SNPs yields 100%.

We constructed 13 scenarios with different proportions k of genetic variance attributed to the QTLs, with 10 independent datasets created for each (Table 1). For the SNPs randomly selected as QTLs and polygenics SNPs, the corresponding effect β_i for selected SNP i was set as follows:

$$\beta_i = \begin{cases} \frac{1}{2} u_i \sqrt{\frac{10^{-4} \sigma_g^2}{2MAF_i(1-MAF_i)}} & \text{if SNP}_i \text{ is polygenic} \\ \frac{1}{2} u_i \sqrt{\frac{k\sigma_g^2}{2MAF_i(1-MAF_i)}} & \text{if SNP}_i \text{ is a QTL} \end{cases},$$

where u_i was drawn from a discrete Uniform $\{-1, 1\}$ distribution to allow non-null effects to take on positive or negative values. For unselected SNPs (i.e., null SNPs), β_i was set to 0. We varied the proportion of genetic variance attributed to each QTL between $k = 0.725\%$ and 5% , with a greater density of values evaluated between 0.725% and 2% ; we focused in particular on this range as it corresponds to more plausible QTL sizes and facilitated a study of the sensitivity of BayesR to small changes. For each value of k , the same $n_{\text{QTL}} = 5$ QTLs were used across scenarios, but the number (and thus the subset) of polygenic SNPs used varied (see Table 1). As the same 5 QTLs were simulated across scenarios for each of the 10 independent datasets, a total of 50 QTLs was considered. Finally, each scenario was run for four different levels of heritability $h^2 = \{0.1, 0.3, 0.5, 0.8\}$, and we evaluated the performance of BayesR for two alternatives: (1) using genotype data that excludes the 5 known QTLs, resembling a classic 50k genotyping array ("50k data"); and (2) using genotype data that includes the 5 known QTLs, which mimics a custom 50k genotyping array ("50k custom data"). In total, this corresponds to $13 \times 10 \times 4 \times 2 = 1040$ simulated datasets.

Statistical Analysis

BayesR genomic prediction model The models of the Bayesian alphabet are all based on the linear model in Equation (1). BayesR assumes that SNP effects β_i follow a four-component normal mixture, making it well-aligned to our simulations (for which SNPs fall into null, weak, and strong classes). The effect of SNP i is assumed to be distributed as

$$\beta_i \sim \pi_1(\beta_i = 0) + \pi_2 N(0, 0.0001\sigma_g^2) + \pi_3 N(0, 0.001\sigma_g^2) + \pi_4 N(0, 0.01\sigma_g^2), \quad (2)$$

where as before, σ_g^2 represents the total additive genetic variance (i.e., the cumulative variance of all SNP effects) and $\pi = (\pi_1, \pi_2, \pi_3, \pi_4)$ the mixing proportions such that $\sum_{i=1}^4 \pi_i = 1$. The mixing proportions π are assumed to follow a Dirichlet prior, $\pi \sim \text{Dirichlet}(\alpha + \gamma)$, with α representing a vector of pseudo-counts and γ the cardinality of each component. In this work, we used a flat Dirichlet distribution, with $\alpha = (1, 1, 1, 1)$, for the prior. As suggested by Moser *et al.* (15), σ_g^2 is assumed to be a random variable following an Inv- χ^2 distribution.

As exact computation of the posterior distribution is intractable for this model, Bayesian inference is performed by obtaining draws of the posterior using a Gibbs sampler; full details of the algorithm can be found in (15) and (18). In practice, at each iteration of the algorithm, SNPs are assigned to one of the four categories, and their effect is subsequently sampled from the full conditional posterior distribution for the corresponding mixture component. Model parameters are then estimated using the posterior mean across iterations, after excluding the burn-in phase and thinning draws. Here, the Gibbs sampler was run for a total of 50,000 iterations, including 20,000 as a burn-in and a thinning rate of 10.

In this work, we used the open source Fortran 90 code described in (15) and available at <https://github.com/syntheke/bayesR>. We made a few modifications to this code, notably adding the posterior variance of estimated SNP effects at each iteration to the output; our modified BayesR code may be found at https://github.com/fmollandin/BayesR_Simulations.

Prediction accuracy for BayesR was quantified using the Pearson correlation between the true phenotypic values (\mathbf{y}) and those estimated using BayesR ($\hat{\mathbf{y}}$) in the validation set.

Statistical criteria for QTL mapping In this section, we present several potential criteria based on BayesR output that can be used for the purpose of QTL mapping. We have sub-divided these criteria into those defined for (1) each SNP individually; (2) neighborhoods, or sliding windows, around each marker; and (3) those used for ranking potential QTLs.

Mapping criteria for individual SNPs

BayesR is unique in the Bayesian alphabet, in that it assigns SNPs to one of four effect size classes at each iteration by weighting according to their likelihood of belonging to each. We thus have access to the posterior frequency with which SNPs were assigned to each class, which can be interpreted as an inclusion probability. We denote the posterior inclusion probability (PIP) of SNP i belonging to class j as $\text{PIP}_i^{(j)}$, such that $\sum_{j=1}^4 \text{PIP}_i^{(j)} = 1 \forall i \in \{1, \dots, p\}$. In the following we interchangeably refer to the null, small, medium, and large classes as $j = 1, 2, 3$, and 4 , respectively. The PIP provides a straightforward method for classifying SNPs as having a null, small, medium, or large effect. We define the maximum a posteriori (MAP) rule for SNP i as

$$\text{MAP}_i = \arg \max_j \text{PIP}_i^{(j)}, \quad (3)$$

implying that SNPs are assigned to their most frequently assigned class. Since SNPs may move frequently from one class to another, the MAP in Equation (3) may not detect SNPs that are predominantly included in the model but move between the three non-null classes. Merging the non-null classes addresses this problem, and

leads to a less stringent criterion, the non-null MAP:

$$\text{MAP}_i^{\text{non-null}} = \begin{cases} 1 & \text{if } \text{PIP}_i^{(1)} < \sum_{j \in \{2,3,4\}} \text{PIP}_i^{(j)} \\ 0 & \text{else} \end{cases} \quad (4)$$

1 Based on this criterion, SNP i is thus included in the model if
 2 $1 - \text{PIP}_i^{(1)} > 0.5$. In this way, all SNPs preferentially assigned to
 3 the null class take on a value of $\text{MAP}_i^{\text{non-null}} = 0$, while those
 4 assigned to any non-null class (small, medium, or large) take on a
 5 value of $\text{MAP}_i^{\text{non-null}} = 1$.

The BayesR model definition explicitly allows for some SNPs to have larger estimated variances than methods such as GBLUP, which tends to shrink the variance of causal marks due to the assumption of a common variance (18). As such, BayesR has the potential for more closely approximating the true variance of QTLs. The posterior variance of SNP i corresponds to

$$V_i = \beta_i^2 X_i^T X_i \quad (5)$$

where X_i represents the i^{th} column of the centered and scaled genotype design matrix. As the SNP effects are computed on the scaled and centered genotype design matrix X , the per-SNP posterior variance can be estimated using

$$\widehat{V}_i = \widehat{\beta}_i^2 X_i^T X_i = \widehat{\beta}_i^2,$$

6 where $\widehat{\beta}_i^2$ corresponds to the posterior mean of β_i^2 , $\widehat{\beta}_i^2 =$
 7 $\frac{1}{N} \sum_{\ell=1}^N \beta_i^{(\ell)2}$, where n is the number of iterations and $\beta_i^{(\ell)2}$ the
 8 value of β_i^2 at iteration ℓ . We indirectly estimated this parameter
 9 as the sum of the posterior variance and squared posterior mean
 10 of each per-SNP effect. We can then estimate a posteriori the pro-
 11 portion of genetic variance of a SNP i as $\widehat{V}_i / \sum_{j=1}^n \widehat{V}_j$.

Neighborhood-based mapping criteria

12 LD represents a preferential association between two alleles and
 13 can have a large impact on how estimated variances are distributed
 14 among SNPs in an LD block. This in turn affects the evaluation
 15 of the variance in the neighborhood of a causal mutation, as well
 16 as the ability to perform QTL mapping using the aforementioned
 17 criteria, for several reasons. First, SNPs in close proximity to a
 18 QTL are likely to be in high LD with it, and thus may erroneously
 19 have their own effects overestimated to the detriment of the QTL's.
 20 The per-SNP criteria defined above risk incorrectly identifying a
 21 QTL as null in such cases. An alternative approach is to define a
 22 neighborhood-based criteria around each marker, thus mapping
 23 QTLs when one or more of its close neighbors is detected. Here,
 24 we define each neighborhood as a sliding window of 15 SNPs
 25 (covering approximately 1Mb) centered around each marker.

Using these neighborhoods, we define the vector of PIPs for a neighborhood centered on SNP i as follows:

$$\text{PIP}_i = (\text{PIP}_i^{(1)}, \dots, \text{PIP}_i^{(4)}) = \text{PIP}_{i'}, \quad \text{with} \quad (6)$$

$$i' = \arg \max_{\ell \in \{i-7, \dots, i, \dots, i+7\}} 1 - \text{PIP}_\ell^{(1)},$$

with the corresponding neighborhood inclusion probability equal to

$$\text{IP}_i = (1 - \text{PIP}_i^{(1)}). \quad (7)$$

The criteria proposed in Equations (3)-(5) can thus be adapted to accommodate neighborhoods as follows:

$$\text{MAP}_i = \text{MAP}_{i'} \quad \text{and} \quad \text{MAP}_i^{\text{non-null}} = \text{MAP}_{i'}^{\text{non-null}}, \quad \text{with} \quad (8)$$

$$i' = \arg \max_{\ell \in \{i-7, \dots, i, \dots, i+7\}} \text{IP}_\ell,$$

where SNP indices are assumed to be ordered according to their physical location. Similarly, the estimated variance of a neighborhood is fixed to the maximal value of its individual markers:

$$V_i = \max_{\ell \in \{i-7, \dots, i, \dots, i+7\}} V_\ell. \quad (9)$$

LD structure raises an additional related problem – in some cases, the BayesR algorithm may alternate assigning different SNPs in an LD block to the large effect class, which has the consequence of diluting variance over a region rather than for a single marker. The window-based criteria in Equations (8)-(9) successfully flag regions where a single SNP sufficiently stands out, but not necessarily those including several diluted effects. In addition, it can be difficult to accurately assess the variance over a region, due to the covariance among SNPs. To provide a neighborhood-level summary of SNP assignments to the four effect classes, we propose the following sliding-window statistic for SNP i , that we will call Weighted Cumulative Inclusion Probability (CIP_i):

$$\text{CIP}_i = \sum_{\ell=i-7}^{i+7} (0 \times \text{PIP}_\ell^{(1)} + 10^{-4} \text{PIP}_\ell^{(2)} + 10^{-3} \text{PIP}_\ell^{(3)} + 10^{-2} \text{PIP}_\ell^{(4)}). \quad (10)$$

Finally, we used the Lewontin D' statistic (19) to quantify the LD between SNPs. Briefly, the LD coefficient D_{AB} between SNPs A and B is defined as $D_{AB} = p_{AB} - p_A p_B$, where p_A , p_B and p_{AB} respectively denote the frequency of allele A in the first locus, allele B in the second, and the frequency of simultaneously having both. D' normalizes D so that $D' = \frac{D}{D_{\max}}$, with

$$D_{\max} = \begin{cases} \max\{-p_A p_B, -(1-p_A)(1-p_B)\}, & \text{if } D < 0 \\ \min\{p_A(1-p_B), (1-p_A)p_B\}, & \text{otherwise.} \end{cases}$$

We will use the maximum value of the LD of a QTL with its neighboring SNPs as a reference for the link disequilibrium in the region.

Criteria ranking for QTL mapping

For the quantitative criteria V_i and CIP_i defined in Equations (9), (5) and (10), we propose the use of rankings for SNP prioritization rather than fixing value thresholds. For QTL mapping based the estimated posterior variance V_i , we focus on the ten SNPs with the highest V_i . As CIP_i represents a sum over 15 SNPs in the neighborhood of SNP i , SNPs adjacent to those that are frequently categorized as non-null are likely to share large values for this criterion. As such, to address this redundancy, we focus on the 150 SNPs with the highest CIP_i value.

Data Availability

The Montbéliarde genotyping data on which simulations are based originate from a private French genomic selection program and were funded by the users (breeding companies and breeders). They are thus proprietary data that cannot be publicly disseminated to the scientific community. All code used to simulate and analyze the data, as well as the scripts to implement BayesR are available on GitHub (https://github.com/fmollandin/BayesR_Simulations).

Prior large class variance	# null	# small	# medium	# large	V_{null}	V_{small}	V_{medium}	V_{large}
0.5%	40783.25	5054.51	300.44	39.80	0	25.55	14.89	9.60
1% (default)	40568.94	5256.72	336.12	16.23	0	26.65	6.91	10.15
2%	40501.21	5307.33	361.08	8.38	0	26.81	17.97	5.45

Table 2 Average (across all simulation scenarios and independent datasets) of the posterior mean cardinality of each BayesR SNP effect class (null, small, medium, large) for three parameterizations of the prior large effect class variance. For a given dataset, each class size (#) is computed as the posterior mean of the number of SNPs assigned to each class across iterations, and V_j is the posterior estimated cumulative variance of each class j .

1 RESULTS AND DISCUSSION

2 Results

3 In the following, we first investigate the sensitivity of BayesR to parameter specification. We next evaluate the model's performance for phenotype prediction and QTL mapping, based on the statistical criteria defined in the previous section, using simulated data that include a set of $n_{\text{QTL}} = 5$ QTLs, as well as polygenic SNPs and null SNPs with no effect on the phenotype.

9 **Sensitivity of BayesR parameter specification** Although the proportion of additive genetic variance assigned to the small, medium, and large effect classes is typically set to 0.01%, 0.1%, and 1% respectively (see Equation (2) and (12)), these prior parameters can be varied by the user. To evaluate the impact on downstream results, we varied the latter between 0.5%, 1%, and 2% for all scenarios with $h^2 = 0.5$, leaving those of the small and medium effect classes at their default values. Modifying the proportion of genetic variance of the large effect class did not appear to have a strong impact on the validation correlation; nevertheless we have observed differences in correlation among the three prior values that can reach 2.6% and 1% for the 50k and 50k custom data respectively. However, we do note that the posterior mean of the number of SNPs assigned in each class and its associated posterior estimated variance appear to be somewhat affected by this parameterization (Table 2). To assess the impact of the prior specification on per-SNP effect estimates, we calculated the Pearson correlation between the estimated posterior means $\hat{\beta}_i$ across SNPs, simulated scenarios and datasets. Among the three prior specifications, the correlation of estimated SNP effects was between 97.4% and 98.6% for all SNPs.

29 Based on these results, we consider that the prior specification appears to have little practical impact on the performance of BayesR, whether for its predictive performance or for per-SNP effect estimates. For the remainder, we therefore use the default prior specification for proportion of genetic variance in each effect class.

35 **Predictive power of BayesR in varied simulation settings** We next sought to investigate the predictive power of BayesR across simulation scenarios, varying the contribution of QTLs to the additive genetic variance (which we refer to as scenarios below), heritability, and use of 50k or 50k custom genotype data.

40 The mean validation correlation (over the ten independent datasets simulated for each) for each simulation scenario illustrates the expected drop in prediction quality for decreasing heritabilities, whether 50k or 50k custom data are used (1). For the former, the mean (\pm sd) validation correlation across scenarios is 0.125 (± 0.048), 0.301 (± 0.057), 0.447 (± 0.058) and 0.650 (± 0.049) for $h^2 = \{0.1, 0.3, 0.5, 0.8\}$. For the latter, the inclusion of the true

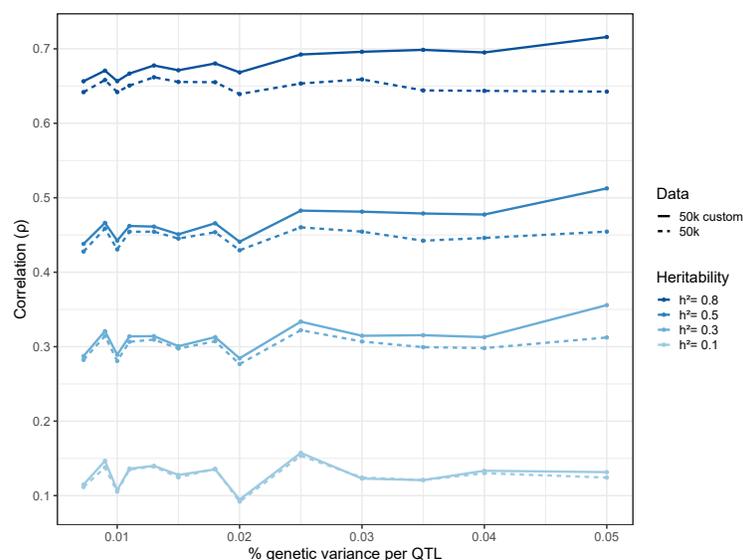


Figure 1 BayesR predictive performance across simulation settings. For each setting (h^2 and percentage of genetic variance assigned to each QTL), points represent mean validation correlations across 10 independent datasets. Heritability values are represented by dark to light blue ($h^2 = 0.8$ to 0.1), and solid and dotted lines represent results for the 50k and 50k custom datasets, respectively.

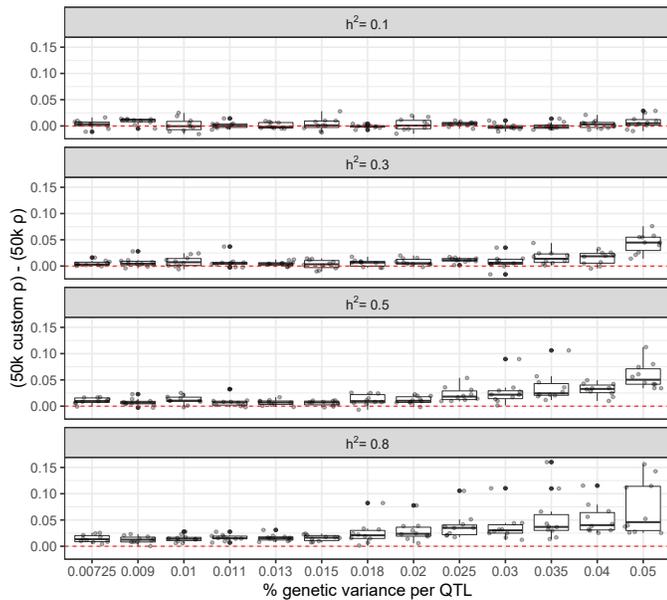


Figure 2 Difference in BayesR predictive performance for the 50k versus 50k custom genotypes across simulation settings. Each panel from top to bottom represents a given heritability ($h^2 = 0.1$ to 0.8), and boxplots represent the distribution of differences in validation correlation between the 50k and 50k custom datasets for each independent dataset (i.e., for which the same 5 QTLs are simulated). The red dotted line indicates a baseline of 0.

1 QTLs among the genotypes unsurprisingly leads to higher validation
 2 correlations, with mean (\pm sd) values across scenarios equal
 3 to 0.128 (± 0.049), 0.312 (± 0.058), 0.466 (± 0.059) and 0.680 (± 0.046)
 4 for $h^2 = \{0.1, 0.3, 0.5, 0.8\}$.

5 Although the trends of the mean validation correlation are
 6 non-linear as the QTL effects take on an increasing percentage of
 7 genetic variance for both types of data, we do remark an increasing
 8 disparity in performance between the 50k and 50k custom data,
 9 particularly as the heritability itself increases (2). In particular, as
 10 expected the potential gain in including the true causal mutations
 11 among genotypes (as is the case of the 50k custom data) appears to
 12 be particularly strong for moderate to large heritabilities and QTL
 13 effects. For $h^2=0.01$, the average difference in validation correlation
 14 was 0.003 (± 0.009), and in some cases the use of the 50k custom
 15 data actually corresponded to a slightly worse prediction. Similar
 16 results are observed at this level of heritability regardless of the
 17 simulated effect size of the QTLs. However, for $h^2 = \{0.3, 0.5, 0.8\}$,
 18 50k custom data led to a nearly systematic gain in performance: the
 19 average increase in validation correlation was 0.011 (± 0.014), 0.019
 20 (± 0.020) and 0.031 (± 0.030) across QTL effect size scenarios, and
 21 attained maximum values of 0.076, 0.112, and 0.160 respectively.
 22 For a given heritability, Figure 2 also shows marked improvements
 23 in prediction when including QTLs simulated with large shares of
 24 additive genetic variance.

25 **QTL mapping using BayesR** A natural first tool to investigate for
 26 QTL mapping is the neighborhood PIP defined in Equation (6). We
 27 focus on the behavior of the neighborhood PIPs for the true QTLs
 28 across scenarios (3), averaging over the 50 QTLs available for each
 29 (5 QTLs \times 10 independent datasets); note that as this is a window-
 30 based measure, this measure can be computed for the true QTLs

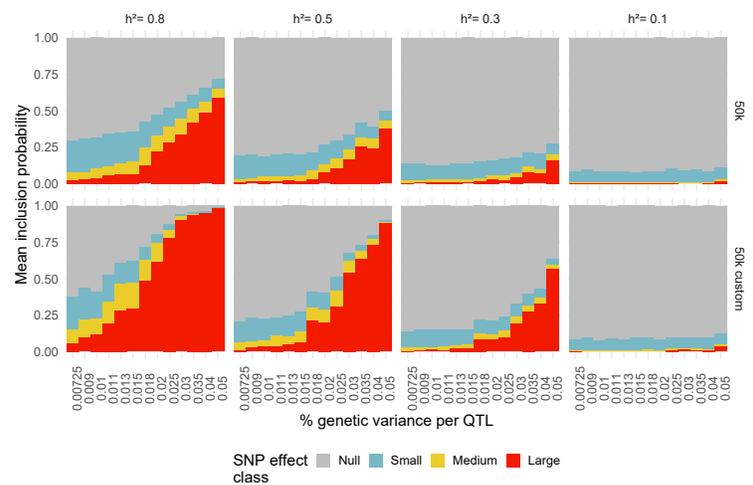


Figure 3 Neighborhood posterior inclusion probabilities across simulation settings. Panels represent combinations of heritability (columns; $h^2 = 0.8$ to 0.1) and type of data used (rows; 50k or 50k custom). Bars represent average (across 5 QTLs \times 10 independent datasets) neighborhood PIP values for the four BayesR effect size classes: null (grey), small (blue), medium (yellow), and large (red).

31 whether the 50k or 50k custom data are used. As shown in 3, the
 32 allotment of true QTL neighborhoods to effect classes varies widely
 33 across heritabilities, proportion of genetic variance for each QTL,
 34 and type of data used. Globally, assigning QTL neighborhoods to
 35 non-null effect classes, particularly the large effect class, is more
 36 frequent for larger heritabilities and simulated QTL effect sizes,
 37 as well as for 50k custom compared to 50k data. However, this
 38 difference disappears for small heritabilities; when $h^2 = 0.1$, the
 39 average (\pm sd) neighborhood PIP for the null class across scenarios
 40 is 0.91 (± 0.009) and 0.90 (± 0.013) for the 50k and 50k custom data,
 41 respectively. Across scenarios, we observe a similar usage of the
 42 small effect class, with an average corresponding neighborhood
 43 PIP of 0.08 (± 0.007) regardless of the genotyping data used. When
 44 $h^2 = \{0.3, 0.5, 0.8\}$, as the simulated share of genetic variances
 45 for QTLs increases for both the 50k and 50k custom data, the null
 46 neighborhood PIP decreases and the large-effect neighborhood
 47 PIP increases. Across all simulated datasets and scenarios, the
 48 average (\pm sd) small- and medium-effect neighborhood PIPs are
 49 0.117 (± 0.053) and 0.058 (± 0.040) respectively, illustrating that
 50 these two classes appear to be less often filled compared to the null
 51 and large classes (although all four classes do appear to be used
 52 outside of the lowest heritability setting).

53 The neighborhood PIP results provide a preview of how QTLs
 54 are grouped into non-null effect classes according to the neigh-
 55 borhood MAP rule (Equation (8); 4). In all simulation settings,
 56 no QTL neighborhoods were assigned to the small effect class
 57 using this criterion. When $h^2 = 0.1$, without surprise, all QTLs
 58 were classified as null. For $h^2 = 0.5$, a very small number of QTL
 59 neighborhoods were assigned to the medium effect class for the
 60 50k data; increasing to $h^2 = 0.8$ led to a larger number moving
 61 to this class for both the 50k and 50k custom data. When not as-
 62 signed to the null class, it was much more common to attribute
 63 QTL neighborhoods to the large effect class; the number of cor-
 64 rectly identified QTL neighborhoods increased with the simulated
 65 effect size and/or heritability, as well as when the causal markers
 66 were included among the genotypes; what's more, these gains tend

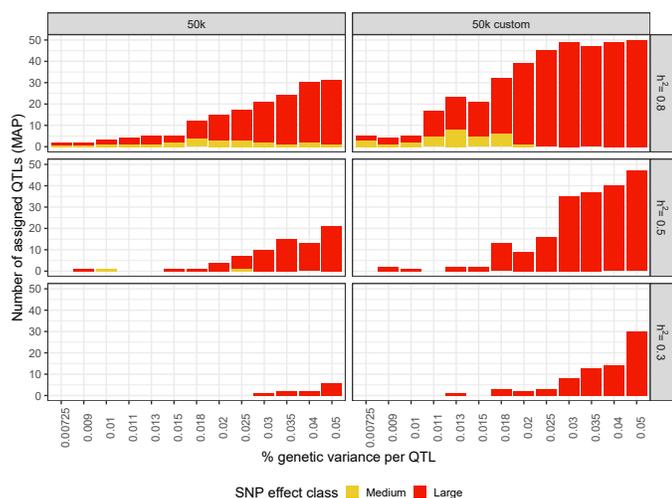


Figure 4 Neighborhood MAP rule for QTL mapping across simulation settings. Number of true QTL windows (out of 5 QTLs \times 10 independent datasets simulated for each scenario, corresponding to a total of 50) correctly assigned to the medium (yellow) and large (red) effect size class using the neighborhood MAP rule. Panels represent data type (columns; 50k and 50k custom) and heritability (rows; $h^2 = 0.8$ to 0.1). The small effect class is not represented because it was empty across all simulation configurations.

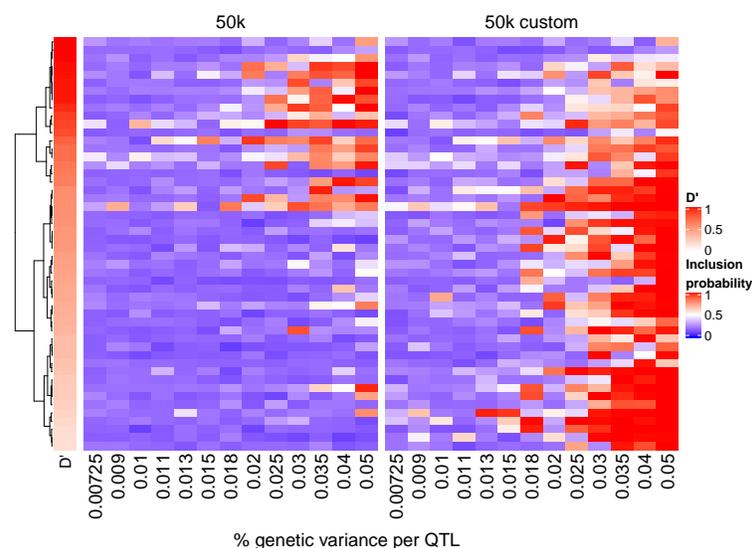


Figure 5 QTL window mapping using the neighborhood inclusion probability across different effect sizes and LD strengths for $h^2 = 0.5$. Neighborhood inclusion probabilities 1-PIP_i⁽¹⁾ for each of the 50 simulated QTLs (heatmap rows) for the 50k (left) and 50k custom (right) data across scenarios (heatmap columns). QTLs are sorted in descending order according to their LD, as measured by D' (left annotation, with deeper reds representing larger values). QTL windows that are represented by white to red cells are correctly detected using the neighborhood non-null MAP.

1 to accumulate when taken together. Correctly detecting at least
 2 one QTL window with the MAP rule required the proportion of
 3 genetic variance simulated for each QTL be $k \geq 3\%$ for $h^2 = 0.3$
 4 using the 50k data, increasing to up to 6 QTL windows for larger
 5 simulated effects. A larger heritability of $h^2 = 0.5$ for the same
 6 data required only $k \geq 0.9\%$ to correctly identify at least one QTL
 7 window, which increases to 22 for $k = 5\%$. However, including
 8 the causal markers in the genotype data enabled detection of QTL
 9 windows at $k \geq 1.3\%$ for $h^2 = 0.3$, with up to 30 correctly detected
 10 at $k = 5\%$. In the most favorable scenario, with $h^2 = 0.8$ and 50k
 11 custom data, QTL windows are detected for all values k , and they
 12 are exhaustively assigned to the large effect class for $k = 5\%$.

13 Given these results, it is not surprising that the neighborhood
 14 $MAP^{\text{non-null}}$ in Equation (8) will tend to detect more QTL win-
 15 dows as being non-null. However, it is also useful to consider
 16 the behavior of this criterion while considering the LD blocks spe-
 17 cific to each simulated QTL. In 5, we visualize the neighborhood
 18 inclusion probability IP_i (defined in Equation (7)) for each of the
 19 50 simulated QTL windows across scenarios for $h^2 = 0.5$, illus-
 20 trating the proportion that are correctly included as non-null in
 21 the model (i.e., when the neighborhood inclusion probability $>$
 22 0.5). The $MAP^{\text{non-null}}$ appears to require a minimum LD of 55%
 23 to correctly recover QTL windows using the 50k data. Below this
 24 threshold, a large portion of QTL windows are not detected. Above
 25 this threshold, QTL window detection appears to become feasible
 26 once the simulated per-QTL percentage of genetic variance attains
 27 about $k = 2\%$. In the 50k custom data, QTL window detection
 28 does not however depend on the amount of LD, although we do
 29 note lower inclusion probabilities for QTLs in very high LD with
 30 their neighbors as compared to the 50k data. Similar to the 50k
 31 data, there is an effect size threshold at about $k = 1.8\%$ at which
 32 QTL windows are more frequently detected.

33 Because the same five QTLs are simulated in each independent

34 dataset across effect size scenarios, Figure 5 also allows for their
 35 specific detection to be followed across configurations. Thus, it can
 36 be seen that some QTLs windows are not detected in any of the
 37 scenarios, while others are more easily detected, even for lower
 38 shares of the genetic variance. That said, there are occasionally
 39 discontinuities in detection observed for increasing shares of the
 40 variance (i.e., a QTL window correctly identified for $k = 0.02$ but
 41 not 0.025). With the exception of $h^2 = 0.1$, which had very weak
 42 detection in all scenarios and datasets, we found similar conclu-
 43 sions for $h^2 = 0.3$ and 0.8, with respectively slightly smaller and
 44 larger overall inclusion probabilities than those shown in Figure 5.

45 Beyond the assignment of SNPs to effect classes using the neigh-
 46 borhood PIPs (and corresponding MAP rules), BayesR also pro-
 47 vides posterior estimates of variability at several levels, including
 48 the additive genetic variance $\hat{\sigma}_g$, the cumulative variance for each
 49 of the three non-null effect classes, and the variance of each SNP.
 50 Before discussing the latter (arguably the most pertinent for QTL
 51 mapping), we verify the estimation quality of the additive genetic
 52 variance. In the 50k genotype data, on average (\pm sd) across scen-
 53 arios, $\hat{\sigma}_g$ was 9.06 (± 3.32), 30.85 (± 3.93), 50.12 (± 4.30) and
 54 77.36 (± 4.61) for $h^2 = \{0.1, 0.3, 0.5, 0.8\}$ respectively; the corre-
 55 sponding true value of σ_g for each were 10, 30, 50 and 80. In the
 56 case of the 50k custom data, this same parameter was estimated to
 57 be 9.11 (± 3.27), 31.01 (± 3.97), 50.27 (± 4.32) and 77.54 (\pm
 58 4.49), respectively.

59 Given that the total additive genetic variance appears to be
 60 well-estimated for both types of genotype type, we turn our at-
 61 tention to the posterior variance $\hat{V}_i / \sum_j \hat{V}_j$ of each neighborhood
 62 as defined in Equation (9). We focus in particular on the case
 63 where $h^2 = 0.5$ and proportions of genetic variance per QTL

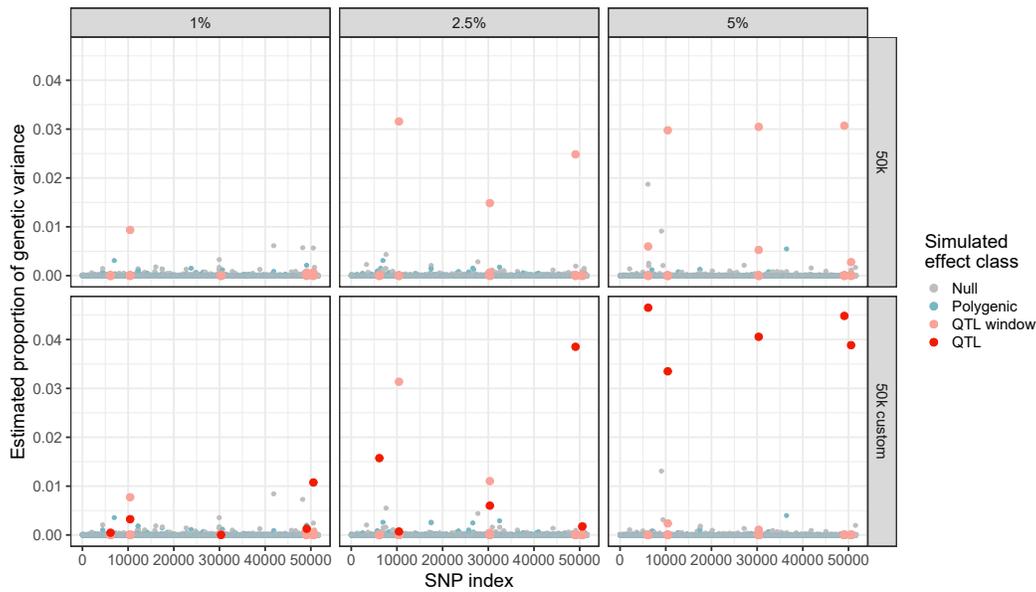


Figure 6 Genome-wide posterior estimate of the proportion of genetic variance per SNP for a single dataset with $h^2 = 0.5$. Posterior estimates of the per-SNP proportion of genetic variance across all $p = 46,178$ SNPs for one of the simulated independent datasets. Panels represent a given simulation setting for percentage of genetic variance per QTL (columns; $k = \{1\%, 2.5\%, 5\%\}$) and data type (rows; 50k versus 50k custom). Points represent individual SNPs, and are colored according to their true effect class (null, polygenic, in the neighborhood of a true QTL, and true QTL). The same five QTLs appear in each panel; true QTLs are only present in the 50k custom data.

1 equal to $k = \{1\%, 2.5\%, 5\%\}$ (6); similar trends were observed
 2 for $h^2 = \{0.3, 0.8\}$. We note that the estimated proportion of ge-
 3 netic variance per SNP window are largely shrunk towards zero,
 4 clearly distinguishing those included in the model. In the 50k
 5 data, certain true QTL windows are clearly prioritized and easily
 6 identifiable. Of the 5 simulated QTLs, we observe one that can
 7 be visually identified for $k = 1\%$, and three for $k = \{2.5\%, 5\%\}$;
 8 more moderated peaks are observed for the remaining QTLs. In
 9 addition, the estimated posterior SNP window variance is about
 10 3%, regardless of the share of variance for the simulated QTLs.
 11 When $k = \{1\%, 2.5\%\}$, the prioritized QTL windows appear to
 12 have estimated variances close to the true simulated values. These
 13 estimates further improve when the 50k custom data are used, and
 14 a larger number of QTLs are clearly prioritized: we note that 2,
 15 4 and 5 QTLs have visibly distinct peaks for $k = \{1\%, 2.5\%, 5\%\}$,
 16 respectively.

17 As a final criterion, we investigate the weighted cumulative
 18 inclusion probability statistic CIP_i defined in Equation (10) as a
 19 way to prioritize neighborhoods where the assignment of SNPs
 20 to non-null classes is somewhat diluted. This statistic tends to up-
 21 weight regions as SNPs in the neighborhood are assigned to non-
 22 null classes (potentially in the place of the primary QTL, which may
 23 be in tight LD with its neighbors). We expect QTL windows already
 24 detected by the neighborhood MAP to similarly have large CIP_i
 25 values; however, it may facilitate the detection of those for which
 26 a cumulative integration of non-null SNPs across the window
 27 provides additional information.

28 To evaluate this point, we compared the QTL mapping perfor-
 29 mance of BayesR using the following three criteria: the neigh-
 30 borhood $MAP_i^{\text{non-null}}$, and the rankings of the neighborhood
 31 V_i (top ten) and neighborhood CIP_i (top 150). We chose to use
 32 $MAP_i^{\text{non-null}}$ here rather than MAP_i as it is less stringent. Across
 33 simulation scenarios and heritabilities, all QTL windows correctly

34 detected by the non-null neighborhood MAP were also identified
 35 by the other two criteria (7). Similarly, all QTL windows correctly
 36 detected by the posterior neighborhood variance V_i ranking were
 37 all also flagged by the CIP_i ranking. The sliding window statistic
 38 thus appears to provide the greatest detection sensitivity, while
 39 the MAP criterion is the most conservative.

40 For all three criteria, the number of detected QTLs increases
 41 with the simulated effect size and heritability, as well as with
 42 their inclusion among the genotypes (50k custom data), with the
 43 exception of the lowest considered heritability, $h^2 = 0.1$. In this
 44 case, no QTL windows are detected with the $MAP^{\text{non-null}}$, and the
 45 number of QTLs identified does not greatly increase for larger QTL
 46 effect sizes. Using the CIP_i rankings, about half of the true QTL
 47 windows can be recovered using the 50k data when $h^2 = 0.8$ in the
 48 50k chip, and similar results are possible with the 50k custom data
 49 for $h^2 = 0.5$. When the true QTLs are excluded from the genotypes,
 50 at most 46 of the 50 true QTL windows can be identified with CIP_i ,
 51 even in ideal circumstances ($h^2 = 0.8$ and $k = 4\%$). However, using
 52 the 50k custom data that include these QTLs allows for universal
 53 detection when $h^2 = 0.5$ and $k = \{3\%, 4\%, 5\%\}$, or $h^2 = 0.8$ for
 54 $k \geq 2.5\%$.

55 Discussion

56 In this work, we evaluated the performance of the BayesR Bayesian
 57 genomic prediction model for prediction quality and QTL map-
 58 ping performance on simulated data under a variety of scenarios,
 59 including varying QTL effect sizes, heritabilities, and the use of
 60 50k versus 50k custom genotype data. Simulated phenotypes were
 61 generated using SNPs from a real set of genotype data in cattle
 62 that were divided into three categories (null, polygenic SNPs, and
 63 QTLs), with variable corresponding shares of the additive genetic
 64 variance. In our study, polygenic SNPs were simulated to have the
 65 same share of genetic additive variance as the default BayesR small

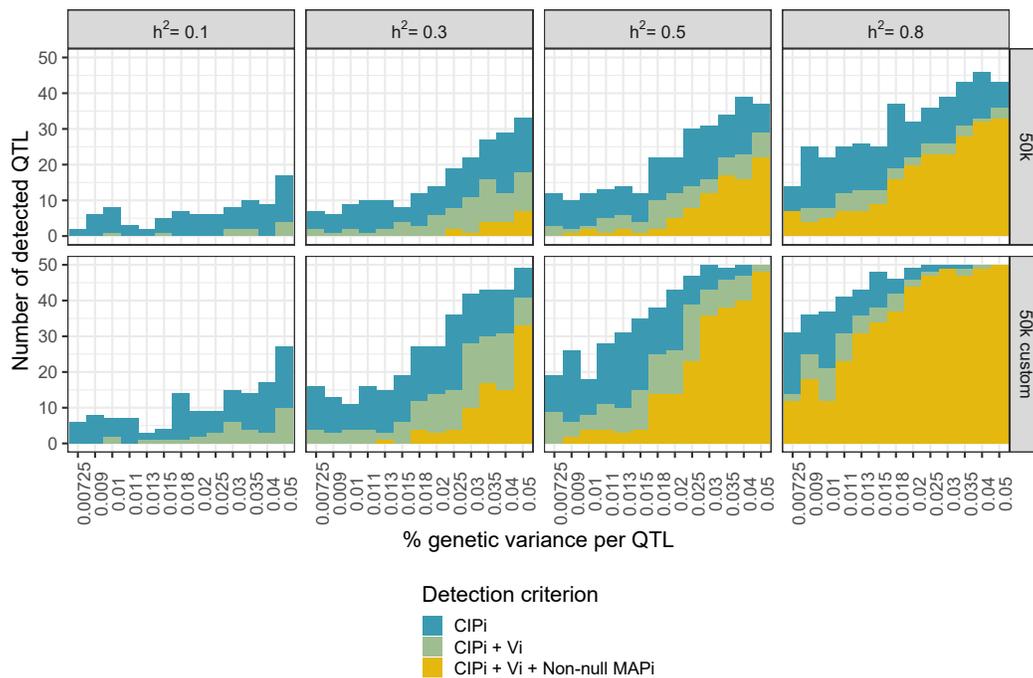


Figure 7 QTL window mapping using three different criteria across simulation settings. Number of true QTL windows (out of 5 QTLs \times 10 independent datasets simulated for each scenario, corresponding to a total of 50) corrected identified using the CIP_i ranking (top 150), V_i (top 10), and $MAP_i^{\text{non-null}}$ neighborhood criteria. Panels represent data type (rows; 50k and 50k custom) and heritability (columns; $h^2 = 0.1$ to 0.8).

1 effect class, i.e. $10^{-4} \times \sigma_g^2$. QTLs were assigned variances ranging
 2 from $7.25 \cdot 10^{-3} \times \sigma_g^2$ to $5 \cdot 10^{-2} \times \sigma_g^2$, constituting an interval that
 3 includes the default prior variance of the BayesR large effect class,
 4 i.e. $10^{-2} \times \sigma_g^2$. These scenarios were simulated at different lev-
 5 els of heritability $h^2 = \{0.1, 0.3, 0.5, 0.8\}$, and we considered both
 6 genotype data that excluded (50k data) or included (50k custom
 7 data) the true simulated QTLs. As the BayesR model definition
 8 includes four different effect size classes (null, small, medium, and
 9 large), it is of particular interest to evaluate how well the model
 10 itself adapts to the underlying genomic architecture of the data.

11 The specific parameterization of BayesR (e.g., number and mag-
 12 nitude of non-null effect classes) can be adapted for different appli-
 13 cations. In this work, we investigated the sensitivity of BayesR re-
 14 sults based on the magnitude of the large effect class, and we found
 15 that the performance of BayesR (predictive power, estimations of
 16 per-SNP effects) was relatively robust. This suggests a limited need
 17 to modifying the priors based on prior biological knowledge.
 18 A more promising approach to integrate such prior knowledge is
 19 the related BayesRC model (13). In the BayesRC approach, SNPs
 20 are divided by the user into two or more non-overlapping subsets,
 21 each of which represents a biologically relevant grouping with
 22 a potentially different proportion of QTLs. For each subset, the
 23 four BayesR SNP effect classes are used, with proportions modeled
 24 using an independent Dirichlet prior (i.e., varying among subsets).
 25 As this flexibility can help prioritize informative SNP subsets that
 26 contain a larger proportion of QTLs, it would be of great interest to
 27 evaluate the impact of the choice of SNP subsets on QTL mapping
 28 with BayesRC, using the criteria we investigated here.

29 With the exception of very low heritability ($h^2 = 0.1$), valida-
 30 tion correlation unsurprisingly increases when QTLs are included
 31 among the genotypes (i.e., the 50k custom data); this increase is

particularly marked for highly heritable phenotypes as well as for 32
 QTLs with large effects. We note that the predictive power of the 33
 BayesR model varied both across simulated scenarios, as well as 34
 within a given scenario, suggesting that the specific position of 35
 simulated QTLs and polygenic SNPs appears to have an influence 36
 on the behavior of BayesR. 37

38 We presented several statistics for QTL mapping and interpreta-
 39 tion using BayesR results, but we note that accurately assessing and
 40 quantifying the importance of a particular genomic region remains
 41 a challenge. One major obstacle is the presence of LD between
 42 SNPs. On one extreme, low LD among neighboring SNPs can
 43 impede the detection of regions if causal mutations are not directly
 44 included among genotypes, while on the other, strong LD blocks
 45 can dilute the signal among adjacent SNPs, leading to alternating
 46 assignments to non-zero effect classes (and subsequently lower
 47 estimated PIPs and variances). While the $MAP_i^{\text{non-null}}$ appears to
 48 be overly conservative for the detection of QTL neighborhoods,
 49 the V_i has the advantage of facilitating an estimation of the pro-
 50 portion of variability corresponding to each QTL neighborhood,
 51 given the overall estimated genetic additive variability. On the
 52 other hand, the CIP_i statistic better takes LD into account by incor-
 53 porating the cumulative importance of an entire region, perhaps
 54 explaining why it can better identify QTL neighborhoods than the
 55 other criteria considered, even under non-optimal conditions (e.g.
 56 $h^2 = 0.1$).

57 There are several limits to our current study that should be
 58 taken into consideration. First, we note that some of our simulation
 59 scenarios could be considered to represent optimal conditions
 60 (e.g., large heritabilities and QTL effect sizes) that would be rare
 61 in real applications. However, studying these extreme scenarios
 62 enables the behavior of the BayesR model to be established in ideal

1 cases. All of our simulations made use of a constant number of
2 individuals in both the training and validation sets, but a future
3 study evaluating the impact of the training population sample size
4 on QTL mapping ability, particularly for cases with low heritability
5 (e.g. $h^2 = 0.1$), could provide insight on this point. Lastly, when
6 sampling SNPs to represent QTLs in our simulations, we chose
7 to limit the choice to those with a MAF > 0.15 , thus excluding
8 those with rare alleles. Although this allowed us to avoid edge
9 cases that would arise with very low MAFs, making it easier to
10 homogenize simulated datasets across different selections of QTLs,
11 this however is an important consideration in QTL mapping.

12 CONCLUSION

13 BayesR is a powerful tool for simultaneously providing accurate
14 phenotypic predictions and mapping causal regions. Our simula-
15 tion results illustrate the flexibility of BayesR for different genomic
16 architectures for all but very low heritabilities ($h^2 = 0.1$) or small
17 QTL effects ($< 1\%$ share of the additive genomic variance). Al-
18 though the four effect size classes (null, small, medium, large)
19 defined in BayesR do not themselves always reflect the true catego-
20 rization of SNPs, they do offer a new approach to understanding
21 and characterizing the genomic architecture underlying a pheno-
22 type. To this end, we presented a variety of statistical criteria that
23 can be used to perform QTL mapping using the output of the
24 BayesR model, including neighborhood-based non-null maximum
25 a posteriori rules, posterior estimated variances, and cumulative
26 inclusion probabilities. We showed that some of the challenges in
27 QTL mapping posed by strong LD blocks could be overcome using
28 the latter criterion, which focuses on the assignment to non-null
29 effect classes of SNPs in an entire neighborhood. By ranking SNPs
30 using this criterion, we demonstrated that QTL windows could
31 more easily be detected, even in simulation scenarios with more
32 challenging conditions.

33 DECLARATIONS

34 Funding

35 This work is part of the GENE-SWitCH project that has received
36 funding from the European Union's Horizon 2020 Research and
37 Innovation Programme under the grant agreement n°817998. This
38 work also benefited from the clustering activities organized with
39 the BovReg project, part of the European Union's Horizon 2020
40 Research and Innovation Programme under the grant agreement
41 n°815668.

42 Acknowledgements

43 The authors thank Mario Calus, Marco Bink, and Bruno Perez
44 for helpful discussions, and Didier Boichard for providing the
45 simulation software used to generate simulated phenotypes from
46 a set of genotype data.

47 ADDITIONAL FILES

48 All code used to simulate and analyze the data, as well as the
49 scripts to implement BayesR are available on GitHub (https://github.com/fmollandin/BayesR_Simulations). The repository is divided into
50 three parts:

- 51 • **Simulations** Fortran source code of the software used to sim-
52 ulate data based on real genotypes. An example of parameters
53 and description of their use are also provided.

- 54 • **bayesR** The modified version of BayesR (available at <https://github.com/syntheke/bayesR>), including recovery of the es-
55 timated per-SNP effects for each iteration, which in turn fa-
56 cilitates the estimation of per-SNP posterior variances. An
57 example of the use of this software is also provided.
- 58 • **codes_R** Partial R scripts used to analyze the BayesR model
59 output and visualize the corresponding results. Scripts to
60 reproduce all figures presented in the article are also included.
61
62

LITERATURE CITED

- 63 [1] Meuwissen THE, Hayes BJ, Goddard ME. Prediction of To-
64 tal Genetic Value Using Genome-Wide Dense Marker Maps.
65 *Genetics*;p. 11.
- 66 [2] Wray NR, Kemper KE, Hayes BJ, Goddard ME, Visscher PM.
67 Complex trait prediction from genome data: contrasting EBV
68 in livestock to PRS in humans: genomic prediction. *Genetics*.
69 2019;211(4):1131–1141.
- 70 [3] Mardis ER. DNA sequencing technologies: 2006–2016. *Nature*
71 *protocols*. 2017;12(2):213.
- 72 [4] Marchini J, Howie B, Myers S, McVean G, Donnelly P. A new
73 multipoint method for genome-wide association studies by
74 imputation of genotypes. *Nature genetics*. 2007;39(7):906–913.
- 75 [5] Pérez-Enciso M, Rincón JC, Legarra A. Sequence-vs. chip-
76 assisted genomic selection: accurate biological information is
77 advised. *Genetics Selection Evolution*. 2015;47(1):1–14.
- 78 [6] Liu A, Lund MS, Boichard D, Karaman E, Fritz S, Aamand
79 GP, et al. Improvement of genomic prediction by integrat-
80 ing additional single nucleotide polymorphisms selected
81 from imputed whole genome sequencing data. *Heredity*.
82 2020;124(1):37–49.
- 83 [7] Brøndum RF, Su G, Janss L, Sahana G, Guldbbrandtsen B,
84 Boichard D, et al. Quantitative trait loci markers derived
85 from whole genome sequence data increases the reliability of
86 genomic prediction. *Journal of dairy science*. 2015;98(6):4107–
87 4116.
- 88 [8] Van den Berg I, Boichard D, Lund MS. Sequence variants
89 selected from a multi-breed GWAS can improve the reliability
90 of genomic predictions in dairy cattle. *Genetics Selection*
91 *Evolution*. 2016;48(1):83.
- 92 [9] Abdollahi-Arpanahi R, Gianola D, Peñagaricano F. Deep
93 learning versus parametric and ensemble methods for ge-
94 nomic prediction of complex phenotypes. *Genetics Selection*
95 *Evolution*. 2020;52(1):1–15.
- 96 [10] Bellot P, de los Campos G, Pérez-Enciso M. Can Deep Learn-
97 ing Improve Genomic Prediction of Complex Human Traits?
98 *Genetics*. 2018 Nov;210(3):809–819. Available from: <http://www.genetics.org/lookup/doi/10.1534/genetics.118.301298>.
- 99 [11] Habier D, Fernando RL, Kizilkaya K, Garrick DJ. Extension
100 of the Bayesian alphabet for genomic selection. *BMC bioin-*
101 *formatics*. 2011;12(1):186.
- 102 [12] Erbe M, Hayes BJ, Matukumalli LK, Goswami S, Bowman PJ,
103 Reich CM, et al. Improving accuracy of genomic predictions
104 within and between dairy cattle breeds with imputed high-
105 density single nucleotide polymorphism panels. *Journal of*
106 *Dairy Science*. 2012 Jul;95(7). Available from: <https://linkinghub.elsevier.com/retrieve/pii/S0022030212003918>.
- 107 [13] MacLeod IM, Bowman PJ, Vander Jagt CJ, Haile-Mariam M,
108 Kemper KE, Chamberlain AJ, et al. Exploiting biological
109 priors and sequence variants enhances QTL discovery and
110 genomic prediction of complex traits. *BMC Genomics*. 2016
111 Dec;17(1):144. Available from: <http://www.biomedcentral.com/1471-2164/17/144>.
- 112
113
114
115

- 1 [14] Sanchez MP, Govignon-Gion A, Ferrand M, Gelé M, Pourchet
2 D, Amigues Y, et al. Whole-genome scan to detect quantitative
3 trait loci associated with milk protein composition in 3 French
4 dairy cattle breeds. *Journal of dairy science*. 2016;99(10):8203–
5 8215.
- 6 [15] Moser G, Lee SH, Hayes BJ, Goddard ME, Wray NR, Viss-
7 cher PM. Simultaneous Discovery, Estimation and Predic-
8 tion Analysis of Complex Traits Using a Bayesian Mixture
9 Model. *PLOS Genetics*. 2015 Apr;11(4):e1004969. Available
10 from: <https://dx.plos.org/10.1371/journal.pgen.1004969>.
- 11 [16] Zhu B, Guo P, Wang Z, Zhang W, Chen Y, Zhang L, et al.
12 Accuracies of genomic prediction for twenty economically
13 important traits in Chinese Simmental beef cattle. *Animal*
14 *genetics*. 2019;50(6):634–643.
- 15 [17] Uemoto Y, Sasaki S, Kojima T, Sugimoto Y, Watanabe T. Im-
16 pact of QTL minor allele frequency on genomic evaluation us-
17 ing real genotype data and simulated phenotypes in Japanese
18 Black cattle. *BMC genetics*. 2015;16(1):134.
- 19 [18] Kemper KE, Reich CM, Bowman PJ, vander Jagt CJ, Cham-
20 berlain AJ, Mason BA, et al. Improved precision of QTL
21 mapping using a nonlinear Bayesian method in a multi-breed
22 population leads to greater accuracy of across-breed genomic
23 predictions. *Genetics Selection Evolution*. 2015 Dec;47(1):29.
24 Available from: <http://www.gsejournal.org/content/47/1/29>.
- 25 [19] Lewontin R. The interaction of selection and linkage. I. Gen-
26 eral considerations; heterotic models. *Genetics*. 1964;49(1):49.