



HAL
open science

RITHMS : An advanced stochastic framework for the simulation of transgenerational hologenomic data

Solène Pety, Mahendra Mariadassou, Ingrid David, Andrea Rau

► **To cite this version:**

Solène Pety, Mahendra Mariadassou, Ingrid David, Andrea Rau. RITHMS : An advanced stochastic framework for the simulation of transgenerational hologenomic data. 2025. hal-04913329

HAL Id: hal-04913329

<https://hal.science/hal-04913329v1>

Preprint submitted on 10 Feb 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

PAPER

RITHMS : An advanced stochastic framework for the simulation of transgenerational hologenomic data

Solène Pety^{1,2*}, Ingrid David³, Andrea Rau¹ and Mahendra Mariadassou²¹Université Paris-Saclay, INRAE, GABI, 78350, Jouy-en-Josas, France, ²Université Paris-Saclay, INRAE, MaIAGE, 78350, Jouy-en-Josas, France and ³Université de Toulouse, INRAE, ENVT, GenPhySE, 31326, Castanet-Tolosan, France

*Corresponding author: solene.pety@inrae.fr

FOR PUBLISHER ONLY Received on Date Month Year; revised on Date Month Year; accepted on Date Month Year

Abstract

A holobiont is made up of a host organism together with its microbiota. In the context of animal breeding, as the holobiont can be viewed as the single unit upon which selection operates, integrating microbiota data into genomic prediction models may be a promising approach to improve predictions of phenotypic and genetic values. Nevertheless, there is a paucity of hologenomic transgenerational data to address this hypothesis, and thus to fill this gap, we propose a new simulation framework. Our approach, an R Implementation of a Transgenerational Hologenomic Model-based Simulator (RITHMS) is an open-source package, builds upon the MoBPS package and incorporates distinctive characteristics of the microbiota, notably vertical and horizontal transmission as well as modulation due to the environment and host genetics. In addition, RITHMS can account for a variety of selection strategies and is adaptable to different genetic architectures. We simulated transgenerational hologenomic data using RITHMS under a wide variety of scenarios, varying heritability, microbiability, and microbiota heritability. We found that simulated data accurately reflected expected characteristics, notably based on microbial diversity metrics, correlation between taxa, modulation of vertical and horizontal transmission, response to environmental effects and the evolution of phenotypic values depending on selection strategy. Our results support the relevance of our simulation framework and illustrate its possible use for building a selection index balancing genetic gain and microbial diversity. RITHMS is an advanced, flexible tool for generating transgenerational hologenomic data that incorporate the complex interplay between genetics, microbiota and environment.

Key words: Holobiont, genotypes, microbiome data, simulation framework

Introduction

It is increasingly understood that the microbiota plays a complex but important role in a variety of biological processes and, more generally, in the construction of host phenotypes. The microbiota is a distinct ecosystem, with its own sources of diversity and modulated by a set of complex factors. In recent decades, the microbiota has been the focus of a growing body of research due to its potential for external modulation [Arnault et al., 2024] or as a selection target [Larzul et al., 2024]. Farm animals live in changing and complex environments, and their microbiota represents a promising way to modulate agroecologically relevant traits, in tandem with their genetics. Indeed, it is now well established that the host genome and its microbiota form a complex organism called a holobiont [Zeng et al., 2015, Bruijning et al., 2022], which corresponds to the ultimate unit on which evolution and selection act [Theis et al., 2016]. Although the microbiota is a complex and dynamic ecosystem, its composition can be explained in part by different modes of transmission. For example, in the first few moments of existence for mammals, maternal contact during delivery and nursing play a crucial role in establishing the initial microbiota through vertical

transmission [Rutayisire et al., 2016, Cortes-Macias et al., 2021]. For non-mammalian vertebrates, such as chickens [Shterzer et al., 2023], the maternal contribution is likely to be considerably weaker than in mammals. In stark contrast to genotypes, a fraction of the microbiota is acquired from the environment through horizontal transmission, and the microbiota continues to evolve throughout a host's life. In addition, both host genes and environmental factors influence the colonization, development, and function of the microbiota, which in turn contributes to host phenotypes.

Given the co-evolution of the host genome and its microbiome under selective pressure, it is of particular interest to integrate microbiota data into genomic prediction models. Such an approach notably offers the potential to improve the prediction of phenotypes and breeding values and has already yielded promising results [Alexandre et al., 2024, Weishaar et al., 2020, Estellé, 2019], although further confirmation is needed in different settings and on larger scales. There is a long tradition of using genetic variants such as single nucleotide polymorphisms (SNP) to estimate breeding values for use in selection schemes [Legarra, 2014], and it is also possible to

construct aggregated selection indices by incorporating traits related to the microbiota, such as taxa diversity. This strategy raises a number of statistical and computational challenges with respect to the simultaneous integration of host genotypes and microbiota. However, benchmarking studies to evaluate predictive hologenomic models require sufficiently large and fully paired transgenerational genomic and microbiota data, notably for the comparison of predicted breeding values to observed offspring phenotypes. Such experimental data are costly to acquire and can be impacted by biases (*e.g.*, fluctuating environmental conditions). Simulation therefore represents an efficient and cost-effective alternative for assessing the relevance of hologenomic prediction strategies.

Several tools for simulating transgenerational genotypes are well known and implemented in user-friendly software such as MoBPS [Pook et al., 2020] or AlphaSim [Faux et al., 2016]. These softwares provide flexible and efficient implementations that allow for a wide range of breeding schemes under a variety of scenarios (*e.g.* heritability) but do not integrate microbiota data. With respect to hologenomic data, other simulation approaches have focused on modeling the structure and dynamics of the microbiota, both for exploring breeding strategies, as well as integrating different types of data to account for complex microbiota-genome interactions [Pérez-Enciso et al., 2021, Wirbel et al., 2022]. Concepts such as transmissibility have also appeared, taking into account the transmissibility of non-genetic information and thus broadening the vision of inclusive heritability [David and Ricard, 2019]. However, none of these hologenomic simulation approaches incorporate a transgenerational aspect. One recent exception is HoloSimR [Casto-Rebollo et al., 2024], a framework that extends AlphaSimR to simulate a co-evolving genome and microbiota under selection. Although it constitutes a solid approach for simulating transgenerational hologenomic data, it does have several limitations: HoloSimR simulations cannot (i) be calibrated based on a real set of hologenomic data, (ii) model the complex heritability of the microbiota, or (iii) include the effects of environmental factors.

Here, we introduce an R Implementation of a Transgenerational Hologenomic Model-based Simulator (RITHMS) to address these limitations. RITHMS is a flexible framework for simulating transgenerational hologenomic data that accounts for the specificities of microbiota transmission and covers the same range of breeding schemes as MoBPS, but under additional and more complex scenarios (heritability, microbiability, microbiota heritability, etc). Real genomic and microbiota data are used to construct a base population from which subsequent generations are derived. This work describes the general framework and strategy used for simulations, and shows that simulated data preserve key characteristics of real data. Finally, it demonstrates the usefulness of transgenerational hologenomic data simulated with RITHMS through a case study on mixed-objective selection, incorporating both phenotype values and microbial diversity.

Material and methods

RITHMS directly incorporates the complexity of transgenerational hologenomic data in several ways (Figure 1): (1) it uses user-provided paired genomic and microbiota data to create a realistic base population from which successive non-overlapping generations are generated, (2) it takes into account the particularities of microbiota transmission (vertical and horizontal) and genetic modulation, (3) it leverages the

functionalities provided in MoBPS [Pook et al., 2020] to define complex genetic architectures and breeding selection steps using indices based on breeding values, microbial descriptors, or a combination of the two, and (4) it facilitates simulations under a variety of scenarios.

RITHMS works with an initialization step of the base population followed by repeated steps (once per generation) of simulation for subsequent generations, summarized in Figure 1 and discussed in greater detail in the following. Key simulation parameters and notations are detailed in Table 1.

Table 1. Table of key parameters and notations for RITHMS

Symbol	Definition	US*
N	Number of individuals in the base population	
n_g	Number of SNPs	
n_b	Number of taxa	
n_{gen}	Number of generations after the base population	✓
n_{ind}	Number of individuals per generation	✓
QTL_y	Number of causative QTL on the phenotype	✓
QTL_o	Number of causative QTL on taxa abundances (per taxon)	✓
OTU_y	Percentage of causative taxa on phenotype	
OTU_g	Percentage of taxa under genetic control	✓
λ	Proportion of vertical transmission	✓
σ_m	Standard deviation value for microbiome noise	✓
$\mathbf{M}^{(t)}$	Taxa abundances	
$\mathbf{G}^{(t)}$	Genotypes encoded as 0,1,2	
$\mathbf{y}^{(t)}$	Phenotype, $\alpha^T \mathbf{G}^{(t)} + \omega^T \mathbf{B}^{(t)} + \epsilon_y^{(t)}$	
$\mathbf{B}^{(t)} = \text{CLR}(\mathbf{M}^{(t)})$	CLR-transformed relative abundance values for taxa of all individuals at generation t ($n_b \times n_{ind}$)	
ω	Taxa effects on phenotype	
α	QTL effects on phenotype	
β	QTL effects on taxa abundances	
σ_β	Standard deviation for non-null cluster- and taxon-specific genetic effects on taxa abundances	✓
$\mathbf{BV}_t^{(t)}$	Total breeding value, $\mathbf{BV}_t^{(t)} = \alpha^T \mathbf{G}^{(t)} + \omega^T \beta \mathbf{G}^{(t)}$	
$\mathbf{BV}_m^{(t)}$	Microbiome-mediated breeding value, $\mathbf{BV}_m^{(t)} = \omega^T \mathbf{B}^{(t)}$	
$\mathbf{BV}_d^{(t)}$	Direct breeding value, $\mathbf{BV}_d^{(t)} = \alpha^T \mathbf{G}^{(t)}$	
h_d^2	Direct heritability, $\text{var}(\alpha^T \mathbf{G}^{(t)}) / \text{var}(\mathbf{y}^{(t)})$	✓
b^2	Microbiability, $\text{var}(\omega^T \mathbf{B}^{(t)}) / \text{var}(\mathbf{y}^{(t)})$	✓
h^2	Total heritability, $\left[\text{var}(\alpha^T \mathbf{G}^{(t)}) + \text{var}(\omega^T \beta \mathbf{G}^{(t)}) \right] / \text{var}(\mathbf{y}^{(t)})$	

*US = User-specified parameters provided as inputs to RITHMS, $^{(t)}$ indicates quantities pertaining to generation t

Formatting of the base population

The base population corresponds to user-provided paired genotype $\mathbf{G}^{(0)}$ and microbiome $\mathbf{M}^{(0)}$ data for N individuals, with respectively n_g SNPs and n_b taxa. For the base population alone (G0), random matings among all individuals are used to generate the following generation (G1) using MoBPS. Exactly n_{ind} individuals are generated according to the user-specified sex ratio (0.5 by default).

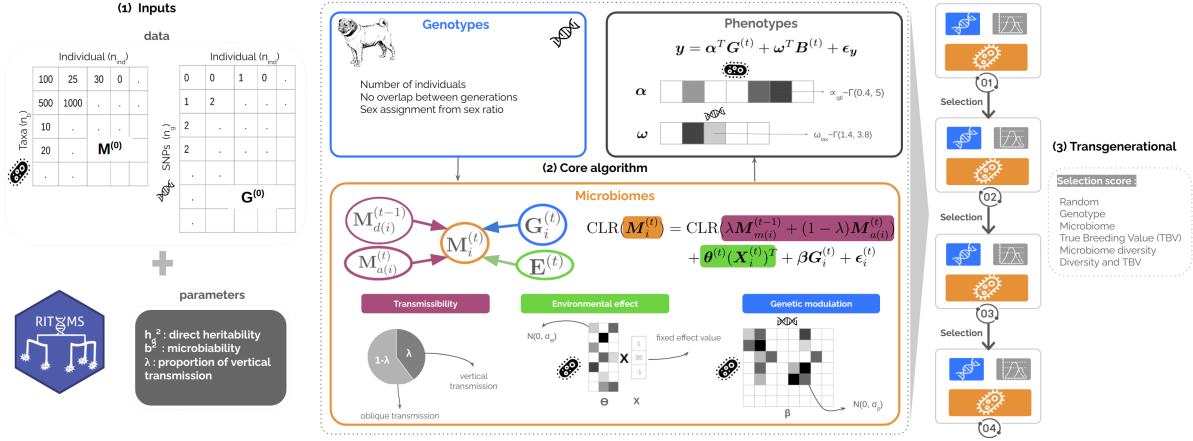


Fig. 1. Overview of RITHMS. (1) User-provided inputs include paired microbiome abundances and genotypes (encoded as 0/1/2) and the following required parameters: direct heritability h_d^2 , microbiability b^2 and λ , which modulates the vertical versus horizontal transmission ratio. (2) For each simulated generation t , the genotypes and pedigree are generated using the MoBPS package [Pook et al., 2020]. Microbiomes are then constructed by first combining maternal and ambient microbiota in proportions λ and $1 - \lambda$ respectively, and subsequently applying genetic and possibly environmental modulation. Genotypes and microbiota are then integrated to simulate the phenotypes of the generation using the recursive model of Pérez-Enciso et al. [2021]. (3) To proceed with the next generation, 30% of the males and 30% of the females are selected, either randomly or based on a selection index chosen by the user.

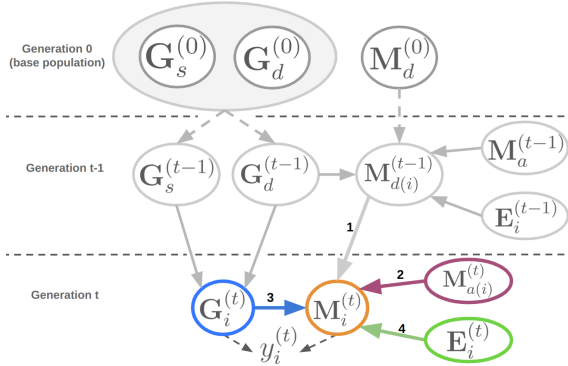


Fig. 2. Schematic illustration of transgenerational hologenomic simulations with RITHMS. The base generation is calibrated on user-provided data for sire genotypes $G_s^{(0)}$, dam genotypes $G_d^{(0)}$, and microbiota data from dams $M_d^{(0)}$. Genotypes are simulated using MoBPS [Pook et al., 2020]. The sources contributing to the variability of taxa abundances of individual i at generation t are as follows: (1) vertical transmission from the individual’s mother $M_d^{(t-1)}$, for example during delivery; (2) horizontal transmission of the ambient microbiota $M_a^{(t)}$; (3) the host selective filter, through which the host’s genotype $G_i^{(t)}$ facilitates the colonization and establishment of certain microorganisms; and (4) individual-specific environmental effects $E_i^{(t)}$, such as diet or treatment effects, modulating the microbiota composition. Phenotypes $y_i^{(t)}$ are simulated according to a linear model, where the microbiome has a direct effect and the genome has both direct and microbiome-mediated effects [Pérez-Enciso et al., 2021].

Genotype data

Genotype data $G^{(0)}$ should be coded as the number of alternative alleles at each variant for each individual (a $N \times n_g$ matrix). Genotypes are provided by the user and can therefore correspond to real data, ensuring realistic linkage disequilibrium (LD) and allelic frequency features in subsequent simulated generations. Sex chromosomes and sample meta-data beyond individual identifiers are not used for the simulations and are ignored if present. Each individual in the population is assigned

as female or male according to a sex ratio parameter (set to 0.5 by default).

Microbiome data

As for the base population genotypes, $M^{(0)}$ can be based on real data. A $N \times n_b$ count matrix is expected, with the same N individuals as those from $G^{(0)}$ as rows and taxa as columns. The provided raw abundances of taxa are used to estimate compositions (*i.e.* vectors of relative abundances) using an empirical Bayes approach to avoid zeroes, and subsequently transformed using the centered log-ratio (CLR) [Aitchison, 1982]. The empirical Bayes approach uses a Dirichlet prior $\mathcal{D}(S\mathbf{p})$ with scale parameter S and mean parameter \mathbf{p} . The latter is set to the population-level composition, estimated as the average of relative abundances across all individuals. In practice, the composition of an individual thus corresponds to a weighted average between its empirical composition and that of the mean population with proportions π and $1 - \pi$ respectively. The default value $\pi = 0.75$ corresponds to a scale parameter S set to a third of the sample total count.

Real hologenomic data used as a base population

To illustrate the functionality of RITHMS, a set of hologenomic data from a single line of $N = 750$ pigs fed a conventional diet [Déru et al., 2020] were used as a base population in this work. Individual pigs were genotyped using a 70K SNP GeneSeek GGP Porcine HD chip, and microbiota composition was analysed using the V3-V4 region of 16S rRNA gene (see Déru et al. [2022] for additional details on data acquisition and processing). We focused here on a subset of the first 5000 SNPs from the chip manifest as well as the 1845 taxa with a prevalence higher than 5% after rarefaction (rarefied depth = 4100 reads).

Simulation of subsequent generations

Simulation of genotypes

We simulate pedigrees and successive generations of genotype data using MoBPS based on the genotype data provided as input. By default, $n_{\text{gen}} = 5$ non-overlapping generations are

simulated, each with $n_{\text{ind}} = 500$ individuals and a sex ratio of 0.5. Regardless of the simulation strategy employed, 30% of females and 30% of males at each generation are chosen to reproduce for the following generation, with selection either performed randomly or based on a user-specific selection score (see Section *Selection*).

Simulation of microbiota

Our simulation framework is based on the idea that the initial composition of an individual's microbiome can be partly inherited from its mother through vertical transmission as well as from the direct environment and is subsequently modulated by the host genotype and environmental factors (Figure 2).

Therefore, we propose the following model for the centered log ratio (CLR) transformed microbiome abundances of individual i at generation t ($t = 1, \dots, n_{\text{gen}}$):

$$\text{CLR}(\mathbf{M}_i^{(t)}) = \text{CLR} \left(\lambda \mathbf{M}_{d(i)}^{(t-1)} + (1 - \lambda) \mathbf{M}_{a(i)}^{(t)} \right) + \boldsymbol{\theta}(\mathbf{X}_i^{(t)})^T + \boldsymbol{\beta} \mathbf{G}_i^{(t)} + \boldsymbol{\epsilon}_i^{(t)}$$

based on the following matrices, with their dimensions :

- $\mathbf{M}_i^{(t)}$: taxa abundances in individual i ($n_b \times 1$)
- λ : proportion of microbiome inherited via vertical transmission from the mother before modulation by selective filtering and random perturbations
- $\mathbf{M}_{d(i)}^{(t-1)}$: taxa abundances of the dam of individual i ($n_b \times 1$)
- $\mathbf{M}_{a(i)}^{(t)}$: Ambient taxa abundances for individual i ($n_b \times 1$)
- $\boldsymbol{\theta}$: environmental effects on taxa abundances ($n_b \times k$)
- $\mathbf{X}_i^{(t)}$: environmental factors for individual i ($1 \times k$)
- $\boldsymbol{\beta}$: multiplicative effect of genotype on taxa abundances ($n_b \times n_g$)
- $\mathbf{G}_i^{(t)}$: genotype of individual i ($n_g \times 1$)
- $\boldsymbol{\epsilon}_i^{(t)} \sim \mathcal{N}(0, \sigma_m^2 \mathbf{I}_{n_b})$: multivariate Gaussian white noise.

Ambient microbiome for each individual Since no herd structure is considered here, we assume that individuals from the same generation live in similar conditions and are therefore exposed to the same sources of microorganisms. We further assume that the ambient microbiome evolves slowly across generations and thus is strongly linked to the average composition in the previous generation, $\overline{\mathbf{M}}^{(t-1)}$. However, each individual will integrate these potential new communities in a different way, leading to the need to include inter-individual variability in the ambient microbiota composition while preserving the structure of the average composition of the previous generation. To introduce inter-individual variability, a random composition is sampled from a Dirichlet prior $\mathbf{M}_{r(i)}^{(t)} \sim \text{Dir}(\eta \overline{\mathbf{M}}^{(t-1)})$, where $\eta > 0$ (set to 25 by default) is the dispersion parameter calibrated to mimic the dispersion in the base population. Here, η was set to 25 based on visual inspection and minimal differentiation of real and simulated compositions according to a PERMANOVA test (results not shown). Similar values of η were found with other datasets [Pérez-Enciso et al., 2021, Chaillou et al., 2015]. Although this sampling results in compositions centered around $\overline{\mathbf{M}}^{(t-1)}$, extremely low abundances, corresponding to very large CLR-transformed values may occur, thus overwhelming any possible modulation by noise or genetics. To regularize this sampling, in particular when $\lambda = 0$ (*i.e.* no vertical transmission), we compute a weighted average between the sampled composition

and the average composition of the previous generation $\overline{\mathbf{M}}^{(t-1)}$:

$$\mathbf{M}_{a(i)}^{(t)} = \pi \mathbf{M}_{r(i)}^{(t)} + (1 - \pi) \overline{\mathbf{M}}^{(t-1)},$$

with the same weights ($\pi, 1 - \pi$) as those used to compute the base population microbiome.

Environmental fixed effects Taxa abundances can be modulated by environmental fixed effects, which may come from different sources and must be modeled accordingly. In particular, some covariates may not impact all taxa, individuals in a generation, or generations in the same way. To incorporate such effects, RITHMS incorporates the term $\boldsymbol{\theta}(\mathbf{X}_i^{(t)})^T$ to indicate the taxa, individuals, and generations for which environmental effects are applied from G1 onwards. In practice, we recommend that nonzero $\boldsymbol{\theta}$ values be drawn from a standard normal distribution.

Genetic modulation It has been shown that some taxa are correlated, regardless of whether or not a taxonomic link exists between them [Coyte and Rakoff-Nahoum, 2019]. It is thus reasonable to assume that the genetic modulation of taxa has a clustered structure so as to mimic existing correlations between taxa. In particular, simulations should yield strongly positive correlations between taxa in the same cluster, and weak or even negative correlations between taxa in different clusters. These clusters are identified from the the rarefied taxa counts of the base population microbiome data $\mathbf{M}^{(0)}$ using hierarchical clustering on Bray-Curtis distances, with 100 clusters by default. To choose the OTU_g taxa (by default 5% of all taxa) under genetic control, we randomly and iteratively select clusters of size 10 to 25 taxa until a threshold of OTU_g taxa is reached.

The challenge then lies in constructing a sparse matrix of QTL effects on taxa CLR abundances, $\boldsymbol{\beta}$. The term $\boldsymbol{\beta} \mathbf{G}^{(t)}$ corresponds to the cumulative effect of the QTL_o causative SNPs on the taxa. Here, by default we set QTL_o to 20% of the total number of SNPs (n_g) divided by the number of clusters under genetic control. To reach a given intra-cluster level of genetic correlation, causative SNPs are sampled randomly for each cluster but common to all taxa within a cluster. The non-null coefficient β_{sg} for genetically modulated taxon s from cluster $c(s)$ with QTL g is then set to $\beta_{sg} = \tilde{\beta}_{c(s)g} + \tilde{\beta}_{sg}$, where $\tilde{\beta}_{c(s)g} \sim \mathcal{N}(0, \sigma_\beta^2)$ and $\tilde{\beta}_{sg} \sim \mathcal{N}(0, \sigma_\beta^2)$. Note that $\tilde{\beta}_{c(s)g}$ depends only on the cluster, ensuring within-cluster correlation, whereas $\tilde{\beta}_{sg}$ is specific to each taxon under genetic control. In this way, as there may be some overlap in QTLs selected for different clusters, both intra- and inter-cluster genetic correlations are induced. The strength of this correlation is mainly limited by the number of clusters and the level of overlap of causative SNPs between clusters. The direction of correlation between taxa in two clusters is given by the sign of $\tilde{\beta}_{c(s)g} \times \tilde{\beta}_{c(s')g}$, summed over the common SNPs between the two clusters. Finally, for each taxon we center $\boldsymbol{\beta} \mathbf{G}^{(t)}$ to ensure both positive and negative genetic modulation, rather than systematic enrichment or depletion within the population.

We have provided a function within RITHMS, `calibrate_gen_effect()`, to help users evaluate the impact of σ_β on the distribution of taxa heritabilities. In practice, it is often reasonable to expect the majority of taxa to have heritabilities on the order of 0.1, with maximum values of no more than 0.5 [Zang et al.].

Quantifying microbiota diversity A variety of metrics exist to quantify the α -diversity (intra-sample diversity) from microbiome data. For a composition $\mathbf{p} = (p_1, \dots, p_{n_b})$, where $\sum_{j=1}^{n_b} p_j = 1$, we consider the Shannon index $H^1(\mathbf{p})$ defined as $H^1(\mathbf{p}) = - \sum_{j=1}^{n_b} p_j \log(p_j)$ with the convention that $0 \log(0) =$

0. When computed directly from sequencing data, this index is based on species counts transformed to relative abundances and thus suffers from potential undersampling of rare species. In order to obtain counts from the relative abundances and mimic this sampling step, n_{ind} multinomial samplings $M(10000, (\mathbf{p}_{i,1}, \dots, \mathbf{p}_{i,n_b}))$ are performed, with $(\mathbf{p}_{i,1}, \dots, \mathbf{p}_{i,n_b})$ the relative abundances of taxa for individual i , equivalent to the cutoffs on sequencing depth used in the dataset analyses described above.

Transgenerational simulation of phenotypes

Phenotypes at generation (t) are simulated as the result of the combined effects of the microbiota and direct genetic effects following the recursive model developed by Pérez-Enciso et al. [2021]:

$$\mathbf{y}^{(t)} = \boldsymbol{\alpha}^T \mathbf{G}^{(t)} + \boldsymbol{\omega}^T \mathbf{B}^{(t)} + \boldsymbol{\epsilon}_y^{(t)}, \quad (1)$$

with:

- $\boldsymbol{\alpha}$ the regression coefficients corresponding to the QTL effects on the phenotype ($1 \times n_g$),
- $\mathbf{G}^{(t)}$ the genotype values of all individuals at generation t ($n_g \times n_{\text{ind}}$),
- $\boldsymbol{\omega}$ the regression coefficients corresponding to taxa effects on the phenotype ($1 \times n_b$),
- $\mathbf{B}^{(t)} = \text{CLR}(\mathbf{M}^{(t)})$, the CLR-transformed relative abundance values for taxa of all individuals at generation t ($n_b \times n_{\text{ind}}$),
- $\boldsymbol{\epsilon}_y^{(t)} \sim \mathcal{N}(0, 1)$, univariate Gaussian noise.

Note that the variance of the Gaussian noise is set to 1 to ensure that changes in mean phenotypic values are expressed in units of standard deviations. In our simulation settings, we assume that all heritable taxa also have an effect on the phenotype; as such, the microbiota effect also includes an indirect genetic effect.

Breeding values and heritability

Under this formulation, we define the Direct Breeding Value as $\mathbf{BV}_d^{(t)} = \boldsymbol{\alpha}^T \mathbf{G}^{(t)}$, the Microbiota-mediated Breeding Value as $\mathbf{BV}_m^{(t)} = \boldsymbol{\omega}^T \mathbf{B}^{(t)}$ and the Total Breeding Value as the expectation of the phenotype given the genotype, $\mathbf{BV}_t^{(t)} = \mathbb{E}[\mathbf{y}^{(t)} | \mathbf{G}^{(t)}] = \boldsymbol{\alpha}^T \mathbf{G}^{(t)} + \boldsymbol{\omega}^T \boldsymbol{\beta} \mathbf{G}^{(t)}$. This $\mathbf{BV}_t^{(t)}$ takes into account both the direct genetic effect ($\boldsymbol{\alpha}^T \mathbf{G}^{(t)}$) due to the transmission of the genotype and the indirect microbiota-mediated ones ($\boldsymbol{\omega}^T \boldsymbol{\beta} \mathbf{G}^{(t)}$), due to the fraction of the microbiota that has an effect on the phenotype and is under genetic control.

From these quantities, it is possible to define a few quantities of interest: (1) the total heritability $h^2 = [\text{var}(\boldsymbol{\alpha}^T \mathbf{G}^{(t)}) + \text{var}(\boldsymbol{\omega}^T \boldsymbol{\beta} \mathbf{G}^{(t)})] / \text{var}(\mathbf{y}^{(t)})$, (2) the direct heritability $h_d^2 = \text{var}(\boldsymbol{\alpha}^T \mathbf{G}^{(t)}) / \text{var}(\mathbf{y}^{(t)})$, and (3) the microbiability $b^2 = \text{var}(\boldsymbol{\omega}^T \mathbf{B}^{(t)}) / \text{var}(\mathbf{y}^{(t)})$.

Parameter calibration The regression vectors $\boldsymbol{\alpha}$ and $\boldsymbol{\omega}$ are fixed across generations and calibrated on the base population. The calibration consists in rescaling $\boldsymbol{\omega}$ based on $\boldsymbol{\alpha}$ in order to reach user-specified values for the direct heritability h_d^2 and the microbiability b^2 . If h_d^2 is set to 0, then all $\boldsymbol{\alpha}$ coefficients are set to zero and the calibration only affects $\boldsymbol{\omega}$. Initial values $\bar{\boldsymbol{\alpha}}$ for the non-zero coefficients of $\boldsymbol{\alpha}$ are sampled from a $\Gamma(0.4, 5)$ distribution and $\bar{\boldsymbol{\omega}}$ for the non-zero coefficients of $\boldsymbol{\omega}$ are sampled from a $\Gamma(1.4, 3.8)$, as done in the Simubiome method [Pérez-Enciso et al., 2021], before rescaling takes place.

Selection To select the individuals that will make up the breeding stock for the next generation, 30% of the males and 30% of the females are selected at each generation. These

fractions can be modified via the parameters `size_selection_F` and `size_selection_M`. If no selection is specified, individuals are chosen at random. Otherwise, a user-specified selection criterion is used to rank individuals, and only a fraction (specified above) of the top performers are retained to reproduce and contribute to the next generation. The available selection criteria are:

- the Direct, Microbiota-mediated, or Total Breeding Values defined above ($\mathbf{BV}_d^{(t)}$, $\mathbf{BV}_m^{(t)}$, $\mathbf{BV}_t^{(t)}$),
- the microbiome diversity, $\delta^{(t)}$, computed as the Shannon diversity,
- a weighted index of microbiome diversity and total breeding value, $w_{\text{div}} \delta^{(t)} + (1 - w_{\text{div}}) \mathbf{BV}_t^{(t)}$, with weight w_{div} set by the user.

Results

In this section, we explore a large number of simulation scenarios to illustrate the capabilities and features of RITHMS. These results were obtained on the dataset described in the *Real hologenomic data used as a base population* section. Unless otherwise specified, all scenarios use the following simulation parameters: $h_d^2 = 0.25$, $b^2 = 0.25$, $\sigma_\beta \times \sqrt{\text{QTL}_o} = 0.3$, $\sigma_m = 0.6$, $n_{\text{ind}} = 500$, $\lambda = 0.5$ and $n_{\text{gen}} = 5$. Other parameters are set to default values as described in the package documentation.

Simulated microbiomes reflect realistic structure

We first evaluate whether the simulated microbiota exhibit expected characteristics. The pairwise correlation matrix of simulated abundances (Figure 3A) shows that RITHMS successfully produces both strong intra-cluster genetic correlations as well as more modest inter-cluster anti-correlations, thanks to the set of partially overlapping QTLs between clusters. Likewise, increasing QTL effect sizes on taxa increases the heritability of taxa abundances, as expected (Figure 3B). The density plots are produced by `calibrate_gen_effect()` and are intended to guide the user in choosing an appropriate effect size to achieve a target distribution of taxa heritabilities. In this setting, a reasonable distribution of taxa heritabilities appears to roughly correspond to a value of $\sigma_\beta \times \sqrt{\text{QTL}_o} = 0.3$. We also confirm the impact of λ in modulating the relative importance of vertical and horizontal transmission (Figure 3C). When $\lambda = 0$, corresponding to no vertical transmission, offspring α -diversity is strongly correlated with that of the ambient microbiota (values averaged over 10 simulated datasets). As λ increases, so does the correlation between maternal and offspring microbiota α -diversity. In contrast, the correlation between paternal and offspring microbiota diversity is low for all values of λ . This is expected, as the sire microbiota does not directly contribute to that of its offspring. Finally, in the absence of selection or environmental filters, the distribution of α -diversity remains stable across generations (Figure 3D), as expected for communities evolving in a neutral framework.

Introduction of sporadic or sustained environmental effects

In breeding and selection programs, it is essential to account for fixed environmental effects, given their strong role in modulating an individual's phenotype. It is therefore important to verify that simulated transgenerational hologenomic data can correctly integrate such factors under a variety of plausible scenarios, such as short-term treatments or long-term diet effects. For the

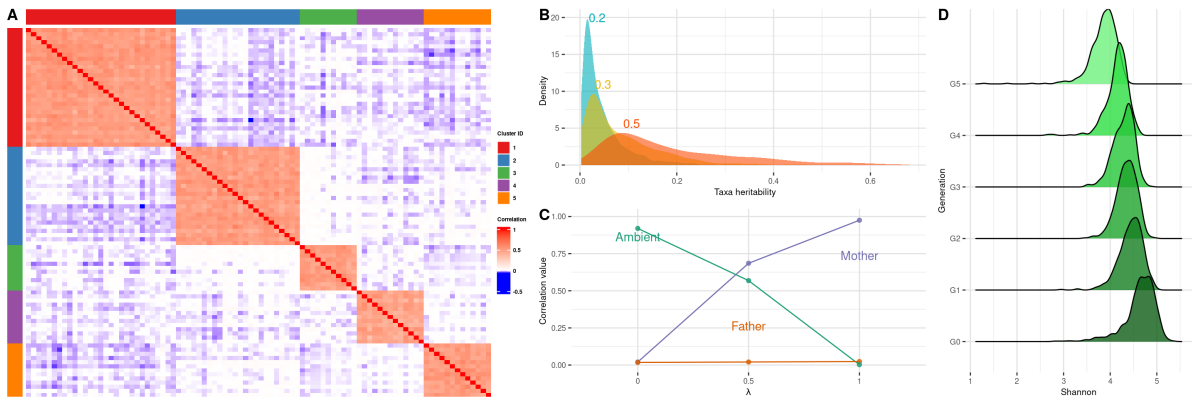


Fig. 3. Key characteristics of microbiota data simulated with RITHMS. (A) Pairwise correlation matrix of taxa abundances. Abundances were simulated assuming all taxa are under genetic control and distributed in five clusters (shown with color bars in the margins). Taxa are sorted based on the cluster they belong to. (B) Density plot of the distribution of taxa heritability for increasing genetic effect sizes ($\sigma_\beta \times \sqrt{QTL_\alpha}$), shown above each curve. (C) Correlation between offspring α -diversity (from G2) and that of its mother (purple), father (orange) or ambient microbiota (green) for increasing values of λ . Correlations are computed from a population of 500 offsprings and averaged over 10 repetitions. (D) Density plots of the distribution of α -diversity values in the base population (G0) and five consecutive generations (G1 to G5), in the absence of selection and environmental filters.

microbiota, as fixed environmental effects can be cumulated with varying effects on each taxa, RITHMS allows users to specify a (potentially sparse) θ matrix, corresponding to the environmental effect sizes on CLR-transformed taxa abundances. To illustrate this, we consider two scenarios introducing either a sporadic (Figure 4A-B) or sustained (Figure 4C-D) environmental effect, as would respectively be the case if a subset of individuals in one generation were administered antibiotics or if individuals in each generation were randomly assigned to different diet groups.

In the first case, we assume that half of the individuals in G1 are administered an antibiotic, provoking significant abundance changes across all taxa. The values for this effect were sampled from a normal distribution $\mathcal{N}(0, 5)$. This one-time environmental effect leads to a strong separation into two groups with very distinct microbiome compositions (Figure 4A) and contrasted α -diversity, as evidenced by the bimodal distribution of α -diversity values in generation G1 (Figure 4B). In the absence of continued antibiotic intake after G1, the lower diversity observed for the antibiotic group is progressively attenuated in the following generations due to random mating, and the bimodality disappears, although the α -diversity is reduced on average compared to the base generation (*e.g.*, when comparing G3 and G0 in Figure 4B). Likewise, the strong group structure in microbiota compositions induced by the treatment progressively disappears in following generations, but the diversity of the overall population shifts towards that of the antibiotic-treated microbiota, suggesting long-lasting changes of the treatment.

In the second case, we assume that individuals from each generation following the base population are randomly assigned to one of two diets, one of which favors abundances in 2 randomly chosen taxa clusters. To simulate a relatively modest effect on the CLR-scale, non-zero values of θ were drawn from a normal distribution with smaller variance than that of the previous case, $\mathcal{N}(0, 2)$. This sustained environmental effect induces a progressive separation of the diet groups that becomes particularly marked at G3 (Figure 4C). As two taxa clusters are preferentially favored in one of the diet groups, with the effect accumulating across generations, we remark the emergence of a group with an increasingly large drop in diversity (Figure 4D).

Impact of genomic, microbiome and hologenomic selection strategies

In the previous sections, we showed that the microbiomes simulated by RITHMS reflect expected characteristics in terms of inter- and intra-cluster genetic correlations among taxa, taxa heritability, vertical or horizontal transmission, as well as microbiome diversity across generations, in the presence or absence of environmental effects. We now turn our attention to phenotypes simulated from the transgenerational hologenomic data under the recursive model in Equation (1). Two critical user-provided parameters for RITHMS simulations are the direct heritability h_d^2 and microbiability b^2 . In the absence of selection, we next sought to verify that the target values are reached and maintained across generations in the case of $h_d^2 = b^2 = 0.25$ (Figure 5A), corresponding to intermediate values and similar to those used in Pérez-Enciso et al. [2021]. h_d^2 and b^2 were computed using the true values of α and ω and the simulated values of $\mathbf{G}^{(t)}$ and $\mathbf{M}^{(t)}$ at each generation. As α and ω are calibrated using the base population to achieve target heritability and microbiability (Section 2.3), it is no surprise that h_d^2 and b^2 are exactly at 0.25 for G0. In subsequent generations, the direct heritability varies only slightly around its target value, and we remark that the observed microbiability tends to be slightly lower than its target value.

Given that the direct heritability and microbiability appear to be reasonably maintained near their target values in the absence of selection, we next evaluate trends in phenotypic improvement as a function of four different selection strategies for varying values of h^2 and b^2 (Figure 5B): selection of 30% of males and 30% of females based on (i) no criterion (random), (ii) the total breeding value ($\mathbf{BV}_i^{(t)}$), (iii) the direct breeding value ($\mathbf{BV}_d^{(t)}$), or (iv) the microbiome-mediated breeding value ($\mathbf{BV}_m^{(t)}$). We observe that the phenotypic change is up to twice as large for higher values of direct heritability and microbiability ($h_d^2 = b^2 = 0.4$) as compared to lower values ($h_d^2 = b^2 = 0.05$). Microbiota selection outperforms the other modes of selection only when microbiability is large compared to the direct heritability ($b^2 = 0.4$ and $h_d^2 = 0.05$). Generally speaking, given the modest contribution of vertical transmission used here ($\lambda = 0.1$, default value), hologenomic selection appears to provide little selection gain compared to genomic selection

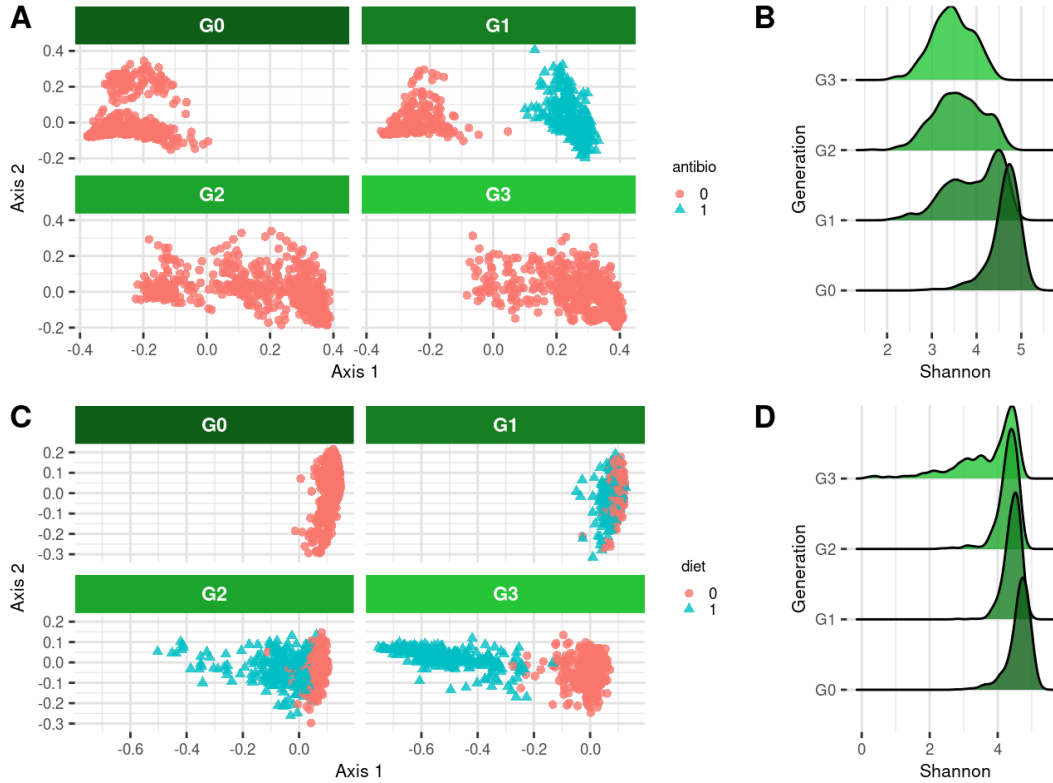


Fig. 4. Simulation of sporadic (top) and sustained (bottom) environmental effects in RITHMS. (A) Multidimensional scaling (MDS) of microbial abundance data (Bray-Curtis distances). Half the individuals at G1 (blue triangles) are subject to a sporadic antibiotic treatment. (B) Density plots of α -diversity values before (G0), during (G1) and after (G2 to G3) sporadic antibiotic treatment. (C) Multidimensional scaling (MDS) of microbial abundance data (Bray-Curtis distances). Starting from G1, half the individuals at each generation (blue triangles) are subject to a diet favoring two clusters of taxa. (D) Density plots of α -diversity values before (G0) and during (G1 to G3) sustained diet intervention.

alone. As an indication, these results were obtained based on a total of 1800 simulated datasets (4 selection modes \times 9 pairs of h_d^2 and b^2 values \times 50 repetitions for each), using the pig hologenomic data described above as a base population, corresponding to 770 minutes of computational time with a maximum memory usage of around 1GB RAM on a laptop with 16 GB RAM (Intel(R) Core(TM) i5-1135G7 CPU @ 2.40GHz \times 8). An implementation for parallelizing repeated simulations is available and demonstrated in the package vignette.

Case study with a mixed selection score

As a final demonstration of the flexibility and usefulness of RITHMS, we consider a practical case study of a complex breeding program with a multi-trait objective: maximizing phenotypic change, based on a quantitative trait of interest $\mathbf{y}^{(t)}$, while preserving microbial α -diversity. One way to achieve this is to use a selection score that combines both objectives into a single value. With access to hologenomic data at each generation, such a score can be constructed as a weighted combination of phenotypic change and diversity. Formally, we define our selection index as $w_{\text{div}} \cdot \delta^{(t)} + (1 - w_{\text{div}}) \cdot \mathbf{BV}_t^{(t)}$ (see the *Selection* section) as a linear combination of the microbial diversity $\delta^{(t)}$ and the total breeding value $\mathbf{BV}_t^{(t)}$, with weight $w_{\text{div}} \in [0, 1]$. Note that $w_{\text{div}} = 0$ corresponds to classic genomic selection. This index is used to identify the 30% of males and 30% of females constituting the breeding stock for the next generation.

Here, we leverage RITHMS to construct a simulation study to identify an optimal weight to achieve gains on both components in a reasonable number of generations. In particular, we simulated data over five generations to evaluate the impact of $w_{\text{div}} \in \{0, 0.1, \dots, 1\}$ on changes in microbial diversity and phenotypic change, with direct heritability $h_d^2 = 0.25$, microbiability $b^2 = 0.25$, vertical transmission $\lambda = 0.5$ and $n_{\text{ind}} = 500$ individuals per generation (Figure 6). Although there is considerable variability among simulated datasets, we remark that there is a tradeoff between mean phenotypic change and microbial diversity (*i.e.*, one comes at the expense of the other), which varies with w_{div} . Larger weights ($w_{\text{div}} = 0.8$ or 0.9) simultaneously achieve phenotypic improvement and increased microbial diversity after five generations. However, for these scenarios, phenotypic change is more modest than for scenarios that increasingly mimic classic genomic selection ($w_{\text{div}} < 0.8$). These results suggest that a value of $w_{\text{div}} = 0.6$ achieves phenotypic change comparable to classic genomic selection in this case study, while drastically limiting the loss of microbial diversity.

Discussion

In this work, we introduced a novel algorithm for simulating transgenerational hologenomic data, implemented in the R package RITHMS. Our tool expands the scope of existing genomic simulation methods [Pook et al., 2020, Gaynor et al.,

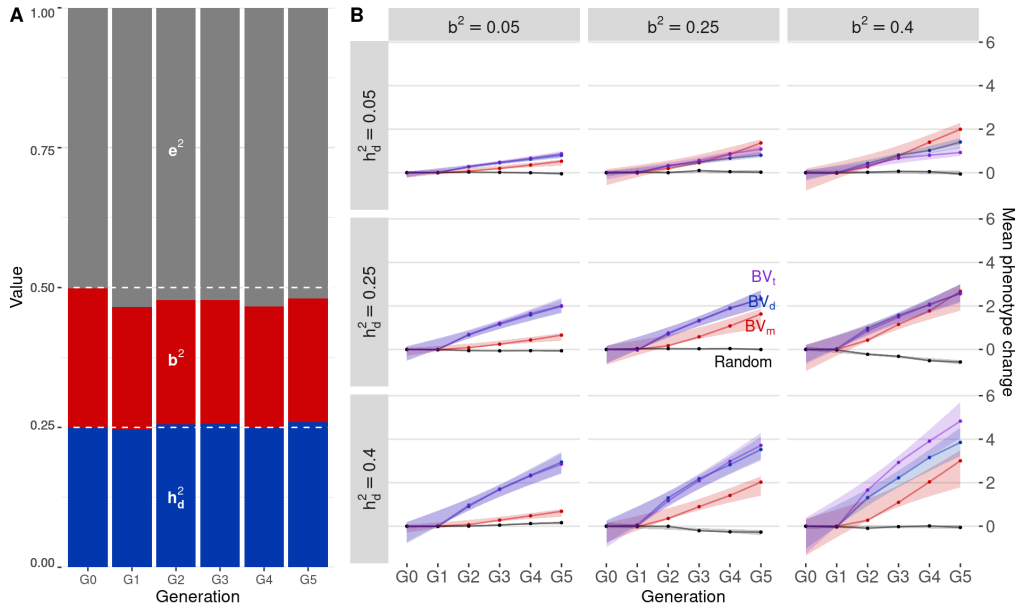


Fig. 5. Direct heritability and microbiability of RITHMS simulations under various selection strategies. (A) Observed direct heritability h_d^2 and microbiability b^2 (averaged over 50 simulated datasets) in a scenario with random selection and target values $h_d^2 = b^2 = 0.25$. (B) Mean phenotypic change across five generations, (averaged over 50 simulated datasets, shaded regions correspond to 95% confidence intervals) with $\lambda = 0.1$, according to various values of direct heritability (rows) and microbiability (columns) and different selection strategies: $BV_d^{(t)}$ (direct breeding values, blue line), $BV_m^{(t)}$ (microbiota breeding values, red line), $BV_t^{(t)}$ (total breeding values, purple line), random selection of parents for the next generation (black line).

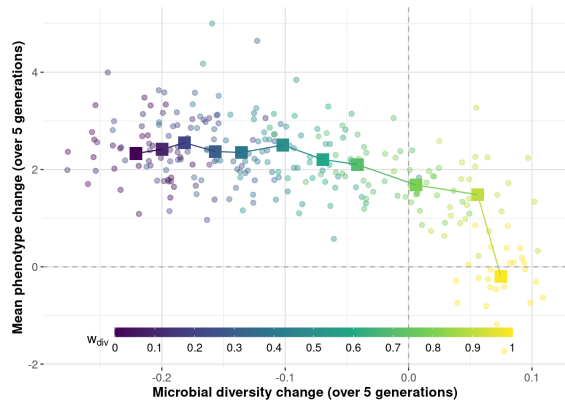


Fig. 6. Simulation-guided exploration of mixed selection index. Mean phenotype and microbial diversity changes from the base population (G0) to G5 as a function of w_{div_0} . The simulation is repeated 25 times for each value of w_{div_0} . Each simulation is shown as semi-transparent dots whereas square dots correspond to the mean computed over the 25 repetitions.

2021] by adding a microbiota compartment and of existing hologenomic simulation methods [Pérez-Enciso et al., 2021] by enabling the simulation of multiple generations. In contrast to the only other transgenerational hologenomic simulator currently available, HoloSimR [Casto-Rebollo et al., 2024], RITHMS uses real data as input, structures the microbiota into taxa clusters and incorporates potential environmental covariates. RITHMS directly accounts for the structure and characteristics of the microbiota as well as its complex transmission mechanisms (from both the mother and the ambient environment, with filters linked to host genetics) and the impact of sporadic or sustained environmental covariates. It is possible to calibrate both (i) the size of genetic effects on the microbiota to obtain

a realistic distribution of taxa heritability and (ii) the direct genetic and microbial effects to achieve target values of direct heritability and microbiability. Complex breeding schemes using the genome, the microbiota or the hologenome combined with different selection scores were used to showcase the flexibility and usefulness of RITHMS. RITHMS is available as an R package, runs on a commercial laptop and is able to generate transgenerational hologenomic data ($n_g = 5000$, $n_b = 2000$ taxa, $n_{ind} = 500$ individuals) for five generations in a few seconds.

Our approach presents several limits and opportunities for future improvements. First, we remark on the slight negative bias we observed between the simulated and target values for microbiability b^2 (Figure 5A, from generation G1 onwards). Since taxa effects on the phenotype ω are calibrated on G0, we hypothesize that this bias originates from a small loss of α -diversity between G0 and G1, as the model cannot reproduce fully the complexity of the base population. Second, our simulation framework is based on a linear model, which has the advantage of being both interpretable and computationally tractable; however, it would be of interest to explore alternatives such as neural networks to introduce non-linearity into RITHMS. Third, our simulated microbiota correspond to snapshots in the lifetime of an animal that are intended for use in predictive models of hologenomic breeding values. However, the microbiota corresponds to a highly dynamic measure that evolves throughout an animal's life, and future work could consider a dynamic model to simulate the microbiota at different time points. Likewise, it would be interesting to extend the RITHMS model to (i) account for microbial interactions with a non-diagonal covariance matrix for the noise component σ_m of the taxa abundances, (ii) allow for the inclusion of more complex environmental effects, and (iii) allow for the use of semi-complete, rather than fully paired, genomic and microbiota data to create the base population, which would enable RITHMS

simulations to be calibrated on a datasets for which some samples lack genomic or microbiota data. Finally, in future work we plan to extend the use of RITHMS to alternative hologenomic datasets, notably for a variety of species and experimental designs, and additional use cases for the evaluation of complex breeding schemes.

Availability

RITHMS is an open-source package available on GitHub (<https://github.com/SolenePety/RITHMS>). A subset of the original dataset is available and can be used to reproduce the figures of the paper in a vignette.

Acknowledgments

This research is supported by the French National Research Agency under the France 2030 program ("ANR-22-PEAE-0006"), which supports future innovation in agricultural science. The authors gratefully acknowledge helpful discussions with Thierry Tribout and Vanille Déru and the breeding companies Axiom and Nucleus for providing animals through France Génétique Porc. Finally, the authors thank the UE3P phenotyping station staff in Le Rheu for animal raising and data recording (<https://doi.org/10.15454/1.5573932732039927E12>).

References

- J. Aitchison. The Statistical Analysis of Compositional Data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 1982.
- P. A. Alexandre, S. T. Rodríguez-Ramilo, N. Mach, and A. Reverter. Combining genomics and semen microbiome increases the accuracy of predicting bull prolificacy. *Journal of Animal Breeding and Genetics*, 2024.
- G. Arnault, C. Marais, A. Préveaux, M. Briand, A.-S. Poisson, A. Sarniguet, M. Barret, and M. Simonin. Seedling microbiota engineering using bacterial synthetic community inoculation on seeds. *FEMS Microbiology Ecology*, 2024.
- M. Bruijning, L. P. Henry, S. K. G. Forsberg, C. J. E. Metcalf, and J. F. Ayroles. Natural selection for imprecise vertical transmission in host-microbiota systems. *Nature Ecology & Evolution*, 2022.
- C. Casto-Rebollo, I. Pocrnic, G. Gorjanc, and N. Ibáñez-Escriche. HoloSimR: A comprehensive framework for simulating breeding programs including the hologenome scenario, 2024.
- S. Chaillou, A. Chaulot-Talmon, H. Caekebeke, M. Cardinal, S. Christieans, C. Denis, M. Hélène Desmonts, X. Dousset, C. Feurer, E. Hamon, J.-J. Joffraud, S. La Carbona, F. Leroi, S. Leroy, S. Lorre, S. Macé, M.-F. Pilet, H. Prévost, M. Rivollier, D. Roux, R. Talon, M. Zagorec, and M.-C. Champomier-Vergès. Origin and ecological selection of core and food-specific bacterial communities associated with meat and seafood spoilage. *The ISME Journal*, 2015.
- E. Cortes-Macías, M. Selma-Royo, I. García-Mantrana, M. Calatayud, S. González, C. Martínez-Costa, and M. C. Collado. Maternal Diet Shapes the Breast Milk Microbiota Composition and Diversity: Impact of Mode of Delivery and Antibiotic Exposure. *The Journal of Nutrition*, 2021.
- K. Z. Coyte and S. Rakoff-Nahoum. Understanding Competition and Cooperation within the Mammalian Gut Microbiome. *Current biology : CB*, 2019.
- I. David and A. Ricard. A Unified Model for Inclusive Inheritance in Livestock Species. *Genetics*, 2019.
- V. Déru, A. Bouquet, C. Hassenfratz, B. Blanchet, C. Carillier-Jacquin, and H. Gilbert. Impact of a high-fibre diet on genetic parameters of production traits in growing pigs. *Animal: An International Journal of Animal Bioscience*, 2020.
- V. Déru, A. Bouquet, O. Zemb, B. Blanchet, M. L. De Almeida, L. Cauquil, C. Carillier-Jacquin, and H. Gilbert. Genetic relationships between efficiency traits and gut microbiota traits in growing pigs being fed with a conventional or a high-fiber diet. *Journal of Animal Science*, 2022.
- J. Estellé. Benefits from the joint analysis of host genomes and metagenomes: Select the holobiont. *Journal of Animal Breeding and Genetics*, 2019.
- A.-M. Faux, G. Gorjanc, R. C. Gaynor, M. Battagin, S. M. Edwards, D. L. Wilson, S. J. Hearne, S. Gonen, and J. M. Hickey. AlphaSim: Software for Breeding Program Simulation. *The Plant Genome*, 2016.
- R. C. Gaynor, G. Gorjanc, and J. M. Hickey. AlphaSimR: An R package for breeding program simulations. *G3 Genes|Genomes|Genetics*, 2021.
- C. Larzul, J. Estellé, M. Borey, F. Blanc, G. Lemonnier, Y. Billon, M.-G. Thiam, B. Quinquis, N. Galleron, D. Jardet, J. Lecardonnel, F. P. Oñate, and C. Rogel-Gaillard. Driving gut microbiota enterotypes through host genetics, 2024.
- A. Legarra. Bases for Genomic Prediction. 2014.
- M. Pérez-Enciso, L. M. Zingaretti, Y. Ramayo-Caldas, and G. de los Campos. Opportunities and limits of combining microbiome and genome data for complex trait prediction. *Genetics Selection Evolution*, 2021.
- T. Pook, M. Schlather, and H. Simianer. MoBPS - Modular Breeding Program Simulator. *G3 (Bethesda, Md.)*, 2020.
- E. Rutayisire, K. Huang, Y. Liu, and F. Tao. The mode of delivery affects the diversity and colonization pattern of the gut microbiota during the first year of infants' life: A systematic review. *BMC gastroenterology*, 2016.
- N. Shterzer, N. Rothschild, Y. Sbehat, J. Dayan, D. Eytan, Z. Uni, and E. Mills. Vertical transmission of gut bacteria in commercial chickens is limited. *Animal Microbiome*, 2023.
- K. R. Theis, N. M. Dheilly, J. L. Klassen, R. M. Brucker, J. F. Baines, T. C. G. Bosch, J. F. Cryan, S. F. Gilbert, C. J. Goodnight, E. A. Lloyd, J. Sapp, P. Vandenkoornhuise, I. Zilber-Rosenberg, E. Rosenberg, and S. R. Bordenstein. Getting the Hologenome Concept Right: An Eco-Evolutionary Framework for Hosts and Their Microbiomes. *mSystems*, 2016.
- R. Weishaar, R. Wellmann, A. Camarinha-Silva, M. Rodehutschord, and J. Bennewitz. Selecting the hologenome to breed for an improved feed efficiency in pigs—A novel selection index. *Journal of Animal Breeding and Genetics*, 2020.
- J. Wirbel, M. Essex, S. K. Forsslund, and G. Zeller. Evaluation of microbiome association models under realistic and confounded conditions, 2022.
- X.-W. Zang, H.-Z. Sun, M.-Y. Xue, Z. Zhang, G. Plastow, T. Yang, L. L. Guan, and J.-X. Liu. Heritable and Nonheritable Rumen Bacteria Are Associated with Different Characters of Lactation Performance of Dairy Cows. *mSystems*.
- Q. Zeng, J. Sukumaran, S. Wu, and A. Rodrigo. Neutral Models of Microbiome Evolution. *PLOS Computational Biology*, 2015.