Identifying Optimal Machine Learning Approaches for Microbiome–Metabolomics Integration with Stable Feature
Selection

Suzette N. Palmer[1-3*], Animesh Mishra[1,4], Shuheng Gan[3], Dajiang Liu[5], Andrew Y. Koh[1,4,6], Xiaowei Zhan[3,6*]

Author information:

[1]Division of Hematology/Oncology, Department of Pediatrics, The University of Texas Southwestern Medical

Center, Dallas, TX 75390, USA

[2]Department of Biomedical Engineering, The University of Texas Southwestern Medical Center, Dallas, TX

75390, USA

[3]Peter O'Donnell Jr. School of Public Health, Quantitative Biomedical Research Center, Center for the Genetics

and Host Defense, The University of Texas Southwestern Medical Center, Dallas, TX 75390, USA

[4]Department of Microbiology, The University of Texas Southwestern Medical Center, Dallas, TX 75390, USA

[5]Department of Public Health Sciences, Pennsylvania State University College of Medicine, Hershey, PA 17033,

USA

[6]Harold C. Simmons Comprehensive Cancer Center, The University of Texas Southwestern Medical Center,

Dallas, TX 75390, USA

Contributions

S.N.P., A.Y.K., and X.Z. planned the research; S.N.P. performed the experiments; S.N.P. analyzed the data; A.M.

and S.G. contributed technical support. S.N.P., A.M., X.Z., D.L., and A.Y.K wrote the paper.

*Corresponding authors

Correspondence to Xiaowei Zhan at xiaowei.zhan@utsouthwestern.edu and/or 214-648-5194.

Keywords: microbiome; metabolomics; multi-omics integration; feature selection; machine learning

Abstract

Microbiome research has been limited by methodological inconsistencies. Taxonomy-based profiling presents challenges such as data sparsity, variable taxonomic resolution, and the reliance on DNA-based profiling, which provides limited functional insight. Multi-omics integration has emerged as a promising approach to link microbiome composition with function. However, the lack of standardized methodologies and inconsistencies in machine learning strategies has hindered reproducibility. Additionally, while machine learning can be used to identify key microbial and metabolic features, the stability of feature selection across models and data types remains underexplored, despite its importance for downstream experimental validation and biomarker discovery. Here, we systematically compare Elastic Net, Random Forest, and XGBoost across five multi-omics integration strategies: Concatenation, Averaged Stacking, Weighted Non-negative Least Squares (NNLS), Lasso Stacking, and Partial Least Squares (PLS), as well as individual 'omics models. We evaluate performance across 588 binary and 735 continuous models using microbiome-derived metabolomics and taxonomic data. Additionally, we assess the impact of feature reduction on model performance and feature selection stability. Among the approaches tested, Random Forest combined with NNLS yielded the highest overall performance across diverse datasets. Tree-based methods also demonstrated consistent feature selection across data types and dimensionalities. These results demonstrate how integration strategies, algorithm selection, data dimensionality, and response type impact both predictive performance and the stability of selected features in multi-omics microbiome modeling.

Introduction

The gut microbiome plays a critical role in host homeostasis and is implicated in the development of various diseases[1,2]. However, microbiome research has largely relied on correlative studies, with limited reproducibility due to methodological discrepancies and data-specific challenges[3–5]. Most microbiome studies focus on taxonomy-based profiling (e.g., 16S rRNA sequencing, metagenomic shotgun sequencing), which introduces challenges such as 1) data sparsity (i.e., a large proportion of zeros representing absent taxa), 2) varying levels of taxonomic resolution, and 3) reliance on DNA-based profiling, which does not necessarily reflect functional microbial activity, such as metabolite or protein production[6,7]. These limitations hinder efforts to move beyond association-based findings toward mechanistic insights.

To address these challenges, multi-omics integration — which combines taxonomic data with functional 'omics (e.g., metabolomics, proteomics, transcriptomics) — has emerged as a promising strategy for uncovering disease-relevant microbial interactions[8–10]. Integrative approaches increase analytical power, allowing researchers to link microbiome composition with functional microbial activity. However, microbiome research methodologies vary widely in data processing, resolution, and integration strategies[11–16] (**Supplemental File 1**). The lack of standardized methodologies can lead to inconsistencies in microbiome-derived findings[17–19]. Differences in machine learning models and integration strategies can yield highly variable results, making it difficult to establish reproducible biomarkers or validate disease associations across studies[10,20,21]. Without a systematic evaluation of these methods, it remains unclear which approaches provide the most biologically meaningful and clinically translatable insights.

To address these gaps, we conducted a multi-faceted analysis to evaluate performance trends and feature selection stability across various multi-omics data integration techniques. Specifically, we analyzed 4 multi-omics datasets with matched microbiome and metabolomics data and assessed integration performance for both binary and continuous outcomes. Our analysis applied Random Forest, XGBoost, and Elastic Net models to both full dimensional and feature reduced datasets. We compared model performance across single 'omics and multi-omics approaches, focusing on two integration strategies: concatenation and stacked generalization. For stacked generalization, we implemented Averaged Stacking, Weighted Non-Negative Least Squares (NNLS)[22], Lasso Stacking (LS)[23], and Partial Least Squares (PLS)[24]. To evaluate the impact of

dimensionality reduction, we filtered and reduced features within each 'omics dataset and compared model performance against the full-dimensional datasets (**Figure 1A**, **Tables 1–2**). Subsequently, we assessed the performance of concatenation, stacked generalization, and single-omics models (**Figure 1B**). Finally, we examined feature selection consistency across machine learning models and integration strategies, comparing stability between single 'omics and concatenated approaches (**Figure 1C**).

Methods

Data Collection and Processing

The data used in this study were originally obtained from Franzosa et al. (2019)[11], Yachida et al. (2019)[14], Erawijantari et al. (2020)[25], and Wang et al. (2020)[26]. Processed taxonomy relative abundance and metabolomics processed data were obtained from Muller et al. (2022)[37] and details can also be found at https://github.com/borenstein-lab/microbiome-metabolome-curated-data/wiki5 and in **Supplemental File 2**.

Data Processing for the Integrated Analysis

Processed data files for metabolomics, metadata, species level relative abundance and genus level relative abundance were further processed for downstream analysis. Taxonomic data at the genus and species levels were concatenated for each dataset, and taxa with an overall abundance across samples that was less than 0.1% were removed. For metabolite differential intensity analysis, Limma 3.64.1[33] was applied, which employs linear modeling and empirical Bayes methods to identify significant differences while adjusting for multiple testing using Benjamini-Hochberg procedure. For taxonomic differential analysis, the Wilcoxon rank-sum test was used to compare taxa abundance across experimental groups (More details for data processing can be found in **Supplemental File 2**).

Overview of Multi-omics Integrative Analysis

The integrative analysis was developed using a custom in-house script that is readily available at https://github.com/suziepalmer10/Multiomics-Integrative-Pipeline/tree/main. The packages tidyverse 2.0.0 and argparse 2.2.3 were used to parse and pre-process the data for this integrative analysis[38]. For continuous outcomes, performance was assessed with Root Mean Squared Error (RMSE), and binary classification was evaluated using the Area Under the Receiver Operating Characteristic Curve (AUROC). For the integrative analysis, AUROC is calculated using pROC 1.18.0 and Kappa is calculated using irr 0.84.1[39,40].

For this analysis, each dataset (as described above, in **Supplemental File 2,** and presented in **Tables**

1–2) was processed through the Integrated Pipeline using Elastic Net, Random Forest, and XGBoost scripts.

The data was split into 80% training and 20% testing, and three repeats of five-fold cross-validation were

performed. To ensure both generalizability and reproducibility, the random seed for each fold was set by

multiplying the repeat index and the fold index (**Supplemental File 2**). Performance was evaluated on validation

and test sets. Grid search with a predefined grid was used for Elastic Net, Random Forest, and XGBoost

(**Supplemental File 2**).

<u>Selection of Machine Learning Methods and Integrative Strategies:</u> Elastic Net, Random Forest, and XGBoost

were chosen because they have been widely used in microbiome research for building predictive models and

identifying the features for experimental validation[41–43]. For multi-omics integration, we adopted feature

concatenation and stacked generalization, two approaches frequently used in this field[11–16]. Collectively,

these algorithms and integration techniques offer interpretable models that can be readily validated through

literature review and experimental follow-up (see **Supplemental File 1** for more details).

<u>Elastic Net Overview and Grid:</u> The Elastic Net algorithm estimates β coefficients to minimize the error term

$L(\beta) = \|Y - X\beta\|^2$, where X represents the data matrix (concatenated, taxonomy, or metabolomics), and Y

represents the response variable. The regularization penalty combines L_1 (Lasso) and L_2 (Ridge) penalties. The

penalty term added to the loss function is $\lambda(\alpha \sum |\beta_j| + (1 - \alpha) \sum \beta_j^2)$, where $\sum |\beta_j|$ represents the L_1 penalty, $\sum \beta_j^2$

represents the L_2 penalty, λ determines the overall regularization strength, and α controls the balance between

L_1 and L_2 regularization. For this analysis, the caret 6.0-93 and glmnet 4.1-6 R packages were used[23,44] (See

Supplemental File 2 for more details).

<u>Random Forest Overview and Grid:</u> The Random Forest algorithm constructs an ensemble of decision trees to

enhance prediction accuracy and reduce overfitting. Each tree is trained on a bootstrapped sample of the data,

and predictions are obtained via averaging (for regression) or majority voting (for classification). To introduce

randomness and reduce feature correlation, Random Forest selects a subset of features at each split. For this

analysis, the caret and randomForest 4.7-1.1 R packages were used[44,45] (See **Supplemental File 2** for more

details).

XGBoost Overview and Grid: XGBoost is an optimized gradient boosting algorithm that iteratively refines weak

learners to minimize the loss function. It incorporates regularization to mitigate overfitting and utilizes efficient

parallel processing for computational efficiency. For this analysis, the caret and xgboost 1.7.8.1 R packages

were used[44,46] (See **Supplemental File 2** for more details).

Data integration using direct concatenation: To integrate taxa and functional data using direct concatenation, we

will horizontally concatenate the two matrices, where A represents taxonomy data and B represents functional

data. Both matrices share the same number of rows (m), corresponding to the samples and disease phenotype

that is used as the response variable. However, these matrices differ in the number of columns, n_A, which

represents the number of taxa and, n_B, which represents the number of functional data (i.e., metabolites).

Mathematically, let $A = \begin{pmatrix} a_{11} & \cdots & a_{1n_A} \\ \vdots & \ddots & \vdots \\ a_{m1} & \cdots & a_{mn_A} \end{pmatrix}$ and $B = \begin{pmatrix} b_{11} & \cdots & b_{1n_B} \\ \vdots & \ddots & \vdots \\ b_{m1} & \cdots & b_{mn_B} \end{pmatrix}$. Concatenation results in C, where $C =$

$[A|B]$. This creates a unified matrix, where $C \in \mathbb{R}^{m \times (n_A + n_B)}$, which enables the integration of the distinct datatypes

while preserving the sample structure. We then use the combined matrix, C, for machine learning model

development using Elastic Net, XGBoost and Random Forest algorithms.

Data integration using averaged stacking of the predictors: Data integration using averaged stacking combines

predictions from individual machine learning models (Random Forest, XGBoost or Elastic Net) for each dataset

(taxonomy and metabolomics) into a single, averaged prediction for each sample. For averaged stacking, let

$P_A \in \mathbb{R}^{m \times 1}$ represent predictions from a machine learning model trained on taxonomy data, A and let $P_B \in \mathbb{R}^{m \times 1}$

represent predictions from a machine learning model trained on functional data. Let m represent the number of

samples and response variables that are shared across both datasets. The averaged stacking prediction, P_{avg},

for each sample is computed as $P_{Avg} = \frac{1}{2}(P_A + P_B)$, where $P_{Avg} \in \mathbb{R}^{m \times 1}$ represents the integration prediction

vector across all samples.

Data integration using weighted non-negative least squares (NNLS): Data integration using weighted nnls

combines predictions from each model by assigning weights to minimize the error between the combined

predictions and the observed values, with the constraint that weights are non-negative. As previously described

above, let $P_A \in \mathbb{R}^{m \times 1}$ and $P_B \in \mathbb{R}^{m \times 1}$, represent predictions from a machine learning model trained on taxa data,

A, and functional data, B. Let $y \in \mathbb{R}^{m \times 1}$ represent response variables and let w_A and w_B represent non-negative

weights assigned to P_A and P_B. The combined prediction $P_{NNLS} = w_A P_A + w_B P_B$, where the weights w_A and w_B

are determined by solving the following optimization problem, where $\min_{w_A, w_B \geq 0} \|y - (w_A P_A + w_B P_B)\|^2$. The nnls 1.5

package was used for non-negative least squares analysis[22].

<u>Data integration using least absolute shrinkage and selection operator (LASSO) stacked generalization:</u> Data

integration using LASSO combines predictions by assigning weights to minimize the error between the combined

prediction and the observed values while enforcing sparsity in the rights through L_1 regularization. As previously

described, let $P_A \in \mathbb{R}^{m \times 1}$ and $P_B \in \mathbb{R}^{m \times 1}$, represent predictions from a machine learning model trained on taxa

data, A, and functional data, B, and let $y \in \mathbb{R}^{m \times 1}$ represent response variables. For LASSO, let w_A and w_B

represent weights assigned to P_A and P_B. The combined prediction, $P_{LASSO} = w_A P_A + W_B P_B$, where the weights

w_A and w_B are determined by solving the following optimization problem using L_1 regularization:

$\min_{w_A, w_B \geq 0} \|y - (w_A P_A + w_B P_B)\|^2 + \lambda(|w_A| + |w_B|)$, where λ is the weight of the L_1 regularization term and controls

solution sparsity. The glmnet 4.1-6 package was using for LASSO[23].

<u>Data integration using partial least squares (PLS):</u> PLS is a regression method that combines the predictors from

separate machine learning models by projecting the predictors into a lower dimensional space to maximize the

covariance between the predictors and the response variables. This integrative method can handle

multicollinearity and captures the shared variation between different datatypes. PLS operates in three stages.

The first stage is latent component extraction, where PLS identifies latent components, T, by projecting the

combined predictor matrix, P, onto a set of weight vectors that maximize the covariance between the latent

components, T, and the response variable, y. Mathematically, $T = PW$, and as previously described, let $P_A \in$

$\mathbb{R}^{m \times 1}$ and $P_B \in \mathbb{R}^{m \times 1}$, represent predictions from a machine learning model trained on taxa data, A, and

functional data, B, and let $y \in \mathbb{R}^{m \times 1}$ represent response variables. Let $P = [P_A | P_B] \in \mathbb{R}^{m \times 2}$, represent the

concatenated matrix of predictors. Let $T \in \mathbb{R}^{m \times k}$ represent the latent variables extracted by P, where k

represents the number of components. Let $W \in \mathbb{R}^{2 \times k}$ represent a weight matrix that defines the linear

combination between P_A and P_B. The next stage of PLS is to create a linear regression of the response variable.

This is accomplished following extraction of latent components, T, and the response variable, y, are then

modeled as a linear function of T. Mathematically, $y = TC + \varepsilon$, where $C \in \mathbb{R}^{k \times 1}$ represent regression coefficients

for T to predict y and ε represents the residual error term. The final stage is to combine the projection and the

regression stages, which can be mathematically represented as $y_{PLS} = PWC$. The pls 2.8-4 package[24] was

used for partial least squares model generation.

Performance Heatmap Generation: Following completion of each machine learning model (Elastic Net, XGBoost

or Random Forest), the performance metrics mean and standard deviation are provided as an output file. The

average AUROC values for binary response variables and the average RMSE values for the continuous

response variables were extracted for each dataset and model using an in-house script. Following creation of

the Heatmap using the Complexheatmap R package[47,48], the top 3 performing models across each dataset

for Elastic Net (if completed), Random Forest and XGBoost. These are represented by the "*", "**", and "***",

and do not include the single 'omics performance metrics. For some of the datasets, the Elastic Net model was

unable to run due to insufficient grid and/or the likely need for further data processing which we decided not to

further optimize for this study.

Performance Pie Charts: The R package ggplot2 3.4.1 was used to create the performance heatmaps for Figures

2-3 and Supplemental Figures 1-2[49]. These metrics were calculated by quantifying the number of models for

each dataset that were part of the top 3 performing models (represented as '*', "**" and "***") (described above

for Performance Heatmap Generation).

Feature Extraction:

Feature importance values were extracted from Elastic Net, XGBoost and Random Forest across all 15 cross-

validation models. For Elastic Net, feature importance is accessed through the magnitude of its regression

coefficients, where larger absolute values reflect greater predictive contributions. For random forest mean

decrease in impurity (MDI) is applied to quantify feature importance and reflects the reduction in impurity

achieved by splitting on a given feature, averages across all decision trees generated. For XGBoost, the Gain

metric is used to access how effectively a feature improves the objective function by reducing prediction error

during tree splits.

Since different metrics were used to quantify feature importance, across Elastic Net, Random Forest and

XGBoost models, we decided to take the absolute value of each variable and perform min-max normalization

(data scaled between 0 and 1) for each model. Mathematically, let i denote the feature ($i = \{1, 2 \dots n\}$), where n

is the total number of features. Let j denote the model ($j = \{1, 2 \dots C\}$), where C represents the number of folds

and repeats from 5-fold cross validation. Let $|S_{i,j}|$ represent the absolute value for the importance score of feature i in model j. Next min-max normalization, $N_{i,j}$, for feature i for model j is calculated by $N_{i,j} = \frac{|S_{i,j}| - min_k(|S_{k,j}|)}{max_k(|S_{k,j}|) - min_k(|S_{k,j}|)}$, where k indexes over all features in model j. I next compute the averaged normalized importance for each feature across all models, $\overline{N}_i = \frac{1}{C}\sum_{j=1}^{C} N_{i,j}$. This allows for direct comparison between features in each model.

For this analysis, the number of top features selected per model was determined by varying feature bin sizes of 1, 5, 10, 20 and 50. For each bin size, only the top n features (where n equals the bin size) were retained to represent the most important features. The features contained in each bin have the greatest importance values and represent the top features that might be explored using peer review and/or by experimentation. Violin plots were generated using ggplot2 to illustrate the overlap and uniqueness of features selected by Elastic Net, Random Forest and XGBoost[49]. The proportions of overlapping and unique features for each dataset are represented as individual circles within a violin plot. Models with fewer than n feature ranked above 0, or those that were not able to finish, were excluded from the analysis.

Results

Creation of integrative analysis pipeline to analyze continuous and binary datasets

We collected and processed four multi-omics datasets consisting of matched metagenomic shotgun sequencing and metabolomics data **(Table 1-2)**[11,14,25,26], and we evaluated 18 continuous and 13 binary matching response variables. We then developed an integrative analysis pipeline to assess the performance of three machine-learning algorithms (Elastic Net, Random Forest, XGBoost) in combination with five integration approaches (Concatenated, Averaged Stacked, Weighted NNLS, Lasso Stacked, and PLS). For concatenation, the metabolomics and metagenomic abundances for each matched sample were appended to form a single long feature vector that served as input for the analysis. The remaining four approaches relied on stacked generalization, where separate base models were first trained on the metabolomics block and on the metagenomics block. The out-of-fold predictions from the two base models were stacked to form two meta-features, which were combined by a meta-learner using all four strategies: a simple mean that assigns equal weight to each prediction (Averaged Stacking), weighted non-negative least square (NNLS) regression, lasso-

penalized regression (Lasso stacking), and partial least squares (PLS). We evaluated each stacking approach on feature reduced and full dimensional datasets to assess how well these models performed under different feature scenarios (**Table 1-2**). Each analysis incorporated a grid search with 3 repeated 5-fold cross-validation to optimize the base models, and performance was assessed using root-mean-squared error (RMSE) for continuous outcomes and area under the receiver-operating-characteristic curve (AUROC) for binary outcomes. Uniform grid parameters were applied across all analyses, and features (or, for stacked models, base model-level importances) were extracted for downstream selection analysis (see **Methods** for more details).

<u>Random Forest with Non-Negative Least Squares Integration Emerges as the Top Performer on Datasets with Continuous Response Variables</u>

We analyzed 364 feature reduced and 371 full dimensional models with continuous response variables across 18 datasets, leveraging 5 integration approaches and 2 individual 'omics datasets (Metabolomics, MSS). To assess model performance, we selected the top 3 models, marked as "*" for the top performer, "**" as the second best, and "***" as the third best performing model for each dataset (**Figure 2**) (See **Methods** for more details).

For feature reduced and full dimensional datasets, Random Forest emerged as the top-performing algorithm. For feature reduced datasets, Random Forest achieved 44.1% of the top results, surpassing XGBoost (31.5%) and Elastic Net (24.1%) (**Figure 2A-B**). For full dimensional datasets, 46.3% of the best performing models were Random Forest, compared to 24.1% Elastic Net and 29.6% XGBoost (**Figure 2C-D**). Among integration methods, Non-Negative Least Squares (NNLS) dominated for both feature reduced and full dimensional datasets (**Figure 2**). For feature reduced datasets, NNLS achieved superior performance in 72.2% of cases, exceeding Averaged Stacked (5.6%), Lasso Stacked (11.1%), PLS (7.4%), and Concatenated (3.7%) (**Figure A-B**). Similarly, in full dimensional datasets, NNLS dominated with 75.9%, while Averaged Stacked (7.4%), Lasso Stacked (9.3%), PLS (7.4%), and Concatenated (0%) fell behind (**Figure 2C-D**).

Interestingly, 9 individual metabolomics models and 1 MSS model outperformed several top integration methods in the features reduced datasets, with 10 metabolomics models and 1 MSS model showing similar trends in full dimensional datasets (**Figure 2A**, **Figure 2C**). Despite significant feature reduction, performance comparisons between feature reduced and full dimensional datasets remained consistent, underscoring the robustness of the Random Forest-NNLS combination for integrating continuous datasets (**Figure 2**).

When examining validation Root Mean Squared Error (RMSE) means, Weighted NNLS demonstrated high performance in both feature reduced (40.7%) and full dimensional (38.9%) datasets, while PLS showed severe overfitting (50.0%) with low testing performance (4.7% for both dataset types) (**Figure 2, Supplemental Figure 1**). Overall, the Random Forest-NNLS combination consistently delivered superior results for datasets with continuous response variables.

Random Forest Maintains Dominance in Binary Outcome Datasets

We analyzed 217 feature reduced and 210 full dimensional models with binary response variables across 13 datasets, applying the same integration approaches and machine learning algorithms as the continuous dataset analysis (**Figure 3**, **Supplemental Figure 2**). Random Forest dominated for both the feature reduced and full dimensional datasets (**Figure 3**). 84.2% of top performing feature reduced models and 72.7% of the top performing full dimensional models were Random Forest (**Figure 3**).

For feature reduced datasets, the integration methods performed comparably among the top models: Lasso Stacked achieved 24.2%, followed by Weighted NNLS and Averaged Stacked each at 21.2%, Concatenated at 18.2%, and PLS at 15.2% (**Figure 3A-B**). In contrast, for full dimensional datasets, Averaged Stacked performed slightly better, making up 30.3% of the top performing models, followed by Lasso Stacked (21.2%), Weighted NNLS (18.2%), PLS (15.2%), and Concatenated (15.2%) (**Figure 3C-D**).

Notably, several single-omics models outperformed the integrative methods in both the feature reduced and full dimensional datasets. Among the feature reduced datasets, 2 metabolomics models and 7 MSS models showed superior performance (**Figure 3A**). In the full dimensional datasets, 4 metabolomics and 4 MSS models outperformed the integrative approaches (**Figure 3C**). Validation AUROC means indicate comparable performance across integration methods, with the exception of PLS (36.1%), which overfit (**Supplemental Figure 2**), similar to the continuous analysis results (**Supplemental Figure 1**). Elastic Net and XGBoost also displayed overfitting, as evidenced by validation AUROC means being significantly higher than testing performance (**Figure 3**, **Supplemental Figure 2**). Overall, the Random Forest algorithm demonstrated robustness across all scenarios, irrespective of integration approach (**Figure 2-3**).

Uncovering Trends in Feature Selection Across Elastic Net, Random Forest and XGBoost

Following model development, we next sought to determine whether our base models (Elastic Net, Random Forest, and XGBoost) identified similar features across concatenated and single-omics (Metabolomics

and Taxa) datasets. Understanding how different models select features across various dataset conditions (feature reduced vs. full dimensional, binary vs. continuous) is crucial for evaluating model consistency and the impact of preprocessing steps on feature selection.

To assess feature selection across models, we first extracted feature importance metrics from each cross-validation step (see **Methods** for more details). For each run, we computed the absolute value of each feature's importance score and applied Min-Max Normalization. This normalization ensures that feature importance values are comparable across models, preventing any single model from dominating due to differences in scale. We then used the Mean Importance to identify the top overall features for each model. For this analysis, we grouped features into bins based on their importance scores within each model, selecting those with the highest feature importance. To ensure rigor, we used predefined bin sizes, which represent the top 1, 5, 10, 20 and 50 features. These discrete bins create a manageable list of high-priority candidates that can be feasibly validated through peer review, or experimentally through *in vivo* and *in vitro* studies. These predefined lists sizes allow us to compare the most influential features across models without the impractical task of vetting each variable across every dataset.

Models were excluded from the analysis if they had fewer than the required number of features with non-zero importance scores, or if Elastic Net failed to complete using the default parameter grid (see **Methods** for more details). Violin plots are used for the following analysis to depict the distribution of the proportions of features selected that are shared and/or unique across Elastic Net, Random Forest and XGBoost for each dataset.

Metabolite Features Selected by Tree-Based Methods Show Strong Concordance, in Contrast to the Distinct Features Selected by Elastic Net

We examined metabolite feature selection patterns from Elastic Net, Random Forest, and XGBoost across feature reduced and full dimensional datasets with binary and continuous response variables (**Figure 4A**). Across all conditions, Elastic Net consistently identified a large proportion of unique features, regardless of dataset dimensionality, response variable type or bin size (1, 5, 10, 20 or 50) (**Figure 4B-E**).

Overlap among the top-ranked features selected by all three models was minimal (**Figure 4B-E**). In nearly all cases, the top feature selected by Elastic Net, Random Forest and XGBoost differed, irrespective of

whether the response variable was binary or continuous, and whether the dataset was full dimensional or feature

reduced. Only one dataset with a continuous response variable and reduced features (**Figure 4B**) and two

datasets with binary response with full dimensional features (**Figure 4E**) showed any overlap in the top ranked

feature across all three models.

Another consistent trend across all conditions were that overlap between Elastic Net and the tree-based

models were limited (**Figure 4**). For feature bins ranging from 5 to 50, overlapping features between Elastic Net

and either, Random Forest or XGBoost accounted for less than 25% of the selected features across all dataset

types. Two exceptions to the trend occurred in the feature-reduced continuous dataset (**Figure 4B**), where these

comparisons showed approximately 40% overlap in the top 5 features for both Elastic Net and Random Forest

and Elastic Net and XGBoost.

Interestingly, in feature reduced datasets with continuous response variables, we observed higher levels

of feature overlap among the three models compared to full dimensional datasets with continuous response

variables (**Figure 4B-C**). This trend, visualized through violin plots of the overlapping distributions, suggests that

dimensionality reduction may increase agreement in feature selection across different modeling approaches for

metabolites with continuous outcomes. This was observed across all conditions, regardless of response variable

type or dimensionality. Violin plots of overlap proportions across bins 1 to 50 also showed a clear pattern of

increased agreement between the two tree-based models (**Figure 4**). These results suggest that Random Forest

and XGBoost tend to select a stable subset of shared features, regardless of dataset structure or outcome type,

in contrast to the increased unique selections made by Elastic Net.

<u>Tree-Models Reveal Partial Concordance in Taxa Selection While Elastic Net Remains Distinct</u>

To assess the consistency and overlap of taxa selected by Elastic Net, Random Forest, and XGBoost,

we analyzed feature selection patterns across full dimensional and feature reduced datasets under binary and

continuous response variables (**Figure 5A**). Across all datasets, there were no instances where the top-ranked

feature was shared among all three models and there were no cases where the top feature was shared between

Elastic Net and Random Forest or between Elastic Net and XGBoost (**Figure 5B-E**). This lack of consensus for

the top feature highlights the distinct feature selection behavior for Elastic Net, XGBoost and Random Forest.

Consistent with these findings, overlap in selected features across bins 5-50 remained low for all model

comparisons with fewer than 25% of features overlapping across all dataset conditions (**Figure 5B-E**).

Patterns of top feature overlap between the tree-based models varied depending on response type and dimensionality. For continuous response variables, Random Forest and XGBoost exhibited greater agreement in their top features when applied to full-dimensional datasets compared to feature reduced datasets (**Figure 5B-E**). This trend was reversed for binary response variables, where tree-based models showed higher concordance in their top selected features in feature reduced datasets than in full dimensional datasets (**Figure 5D-E**).

Violin plot distributions of feature overlap between Random Forest and XGBoost for bins 5-50 maintained a consistent pattern across all conditions (**Figure 5B-E**). The distributions of shared features closely resembled those of features uniquely identified by each model. On average, approximately half of the selected features in these bins were shared between the two tree-based methods, while the other half were distinct to either Random Forest or XGBoost (**Figure 5B-E**). This indicates that, although the tree-based models exhibit some level of agreement, each still captures a large subset of unique taxa features.

In contrast, Elastic Net consistently selected a largely distinct set of features across all datasets and bin sizes. For bins 1-50, more than 75% of the selected features were unique to Elastic Net for the majority of datasets (**Figure 5B-E**). The high proportion of unique selections further underscores the importance of model selection and comparison of features selected by each model.

Distinct Elastic Net Selections Contrast with Consistent Tree-Based Features for Concatenated Models

To assess the consistency and overlap of features selected from concatenated models, which combine both taxa and metabolite data, we evaluated feature selection patterns from Elastic Net, Random Forest and XGBoost across full dimensional and feature reduced datasets, using binary and continuous response variables (**Figure 6A**). The overall trends for continuous response variables for the concatenated models closely mirrored those observed in the taxa-only models (**Figure 5B-C, Figure 6B-C**). Specifically, overlap in the top feature selection across all three models was rare (**Figure 6B-C**). Only one feature reduced dataset showed agreement among all three models for the top ranked feature, and no overlap was observed in full-dimensional datasets or for any of the binary response datasets (**Figure B-E**).

Only one continuous dataset showed overlap between Elastic Net and XGBoost for the top feature selected, while no overlap was observed between Elastic Net and Random Forest across any dataset, regardless of response type or dimensionality (**Figure 6B-E**). In contrast, Random Forest and XGBoost exhibited relatively

high agreement for their top selected features across binary and continuous datasets, with approximately half of

the proportion of features being shared (**Figure 6B-E**). Elastic Net consistently selected distinct top features for

continuous datasets, with the exception of two feature reduced datasets and one full dimensional dataset (**Figure**

6B-C).

When examining the top selected features (bins 5-50, which represent the top 5 through 50 features

selected per dataset), this trend of limited agreement between Elastic Net and tree-based models continued.

The proportion of overlapping features among all three models, as well as between Elastic Net and either tree-

based methods remained low across most datasets. Overlapping proportions exceeding 25% were rare, and the

majority of distributions were centered near zero, further confirming that Elastic Net largely selected features

distinct from those chosen by Random Forest and XGBoost (**Figure 6B-E**). Similar to the taxa-only models,

Random Forest and XGBoost demonstrated greater consistency in their feature selections for the concatenated

datasets (**Figure 5-6**). Across bins 5-50, approximately 50% of the selected features were shared between the

two tree-based model, while the remaining features were largely unique to each method (**Figure 6B-E**). These

results reinforce the pattern that Elastic Net tends to prioritize a separate set of features, while Random Forest

and XGBoost show more concordance in their selections, even when analyzing integrated taxonomic and

metabolomics data.

Dimensionality Reduction Influences Feature Selection Stability in Machine Learning Models

To evaluate the impact of dimensionality reduction on model behavior, we compared feature selection

between full dimensional and feature reduced datasets for each of the machine learning methods, Elastic Net,

Random Forest and XGBoost, across all metabolomics, taxa and concatenated datasets using binary and

continuous response variables (**Figure 7A**). Our goal was to determine whether dimensionality affected feature

selection stability and whether the top features selected remained consistent across the different machine

learning algorithms.

Across all data types and response variables, Elastic Net consistently demonstrated minimal agreement

in top feature selection for the full dimensional and feature reduced models (**Figure 7B-G**). Regardless of

whether the dataset involved metabolomics, taxonomy or concatenated or the response type, only 1-2 datasets

showed overlap in the top feature selected by Elastic Net (**Figure 7B-G**). In contrast, the tree-based models,

with greater emphasis on Random Forest, showed stronger consistency in the top features selected for the

metabolomics datasets (**Figure 7B, Figure 7E**). Approximately half of the Random Forest models identified the same top feature in both the full and feature reduced datasets for both binary and continuous response variables. XGBoost followed a similar trend for metabolomics data, with slightly lower proportions of agreement in top features (**Figure 7B, Figure 7E**).

In contrast, for taxa and concatenated datasets, top feature agreement between full dimensional and reduced models was rare, even for the tree-based methods (**Figure 7C-F**, **Figure 7D-G**). Across these analyses, Random Forest and XGBoost typically shared the same top feature in only 1-4 datasets indicating that dimensionality reduction substantially influences which features are prioritized for these data types (**Figure 7C-F**, **Figure 7D-G**).

When examining bins 5-50, further insights emerged regarding the broader set of selected features. For metabolomics data, regardless of response type, the majority of datasets showed less than 50% overlap in features between the full dimensional and feature reduced models (**Figure 7B, Figure 7E**). Most distributions were centered around zero, with emphasis on Elastic Net, which displayed major differences in selected features. Notably, Random Forest models applied to metabolomics datasets demonstrated relatively even distributions of overlap across bins 5-50, while XGBoost distributions differed based on response type (**Figure 7B, Figure 7E**). For continuous outcomes, XGBoost showed a concentration of overlapping proportions in the 0.0-0.20 range, whereas binary outcomes exhibited a more even spread across all feature bins (**Figure 7B, Figure 7E**).

For taxa data, feature overlap across bins 5-50 was sparse, with the highest density of violin plot distributions concentrated at zero for both binary and continuous response types (**Figure 7C, Figure 7F**). Some exceptions were observed, which include one continuous taxa dataset at bin 5 and two binary taxa datasets at bin 5. These showed complete agreement in selected features between full and reduced models (**Figure 7C, Figure 7F**). Random Forest models developed using taxa data showed a more evenly distributed pattern of overlap, whereas XGBoost had a concentration near zero but still demonstrated higher concordance than Elastic Net across most bins (**Figure 7C**, **Figure 7F**).

Trends observed in the concatenated models closely mirrored these seen in the taxa data (**Figure 7D**, **Figure 7G**). However, distributions for bins 5-50 were generally shifted closer to zero across all three modeling approaches, indicating even lower agreement in selected features between full dimensional and feature reduced models (**Figure 7D**, **Figure 7G**). This suggests that when combining multiple data types, feature selection

becomes increasingly sensitive to dimensionality, especially for Elastic Net, which consistently demonstrated the

least concordance between full and feature reduced models.

<u>Differences in Concatenated and Single 'Omics Feature Selection</u>

To determine whether concatenated models incorporating both metabolomics and taxa data could

recover features identified by single 'omics models, we compared feature selection patterns across Elastic Net,

Random Forest and XGBoost for both full dimensional and feature reduced datasets with continuous and binary

response types (**Figure 8A**, **Figure 8F**). For metabolomics data, Elastic Net models trained on feature reduced

concatenated datasets showed no overlap in features with the corresponding single 'omics metabolomics models

across bins 1-50, regardless of response type (**Figure 8B-E**). In the full dimensional datasets, there was limited

agreement between the concatenated and metabolomics models for continuous outcomes, with most overlap

proportions falling below 0.4 and the highest density occurring near zero (**Figure 8C**). For full dimensional data

with binary outcomes, consensus remained minimal with many datasets showing no overlapping features

between the concatenated and metabolomics models (**Figure 8D-E**).

In contrast, Random Forest and XGBoost models demonstrated a greater ability to identify overlapping

metabolite features between concatenated and metabolomics single 'omics datasets (**Figure 8B-E**). Across top

feature bins 1-50, both models recovered shared features more proportionally, and in several datasets, the same

top features were selected. For feature reduced datasets, however, this agreement diminished, particularly for

continuous response variables (**Figure 8B-C**). The strongest agreement was observed in the full dimensional

models, where overlap between the concatenated and metabolomics datasets were more evenly distributed

across bin sizes (**Figure 8C**, **Figure 8E**). A notable exception was seen in the top 50 metabolite bin for XGBoost,

where the distribution was reduced (**Figure 8C**, **Figure 8E**).

When comparing taxa features between concatenated and single-omics models, we observed a higher

degree of consensus relative to the metabolite comparisons. Elastic Net models, surprisingly, displayed

substantial overlap across all bins, with the majority of feature overlap proportions falling in the 0.75 to 1.0 range

across both full dimensional and feature reduced datasets, regardless of response type (**Figure 8G-J**). An

exception occurred in the top feature comparison for the full dimensional binary dataset, where only one dataset

showed agreement between the concatenated and taxa models (**Figure 8J**).

For Random Forest and XGBoost, the proportion of overlapping taxa features between concatenated and single 'omics datasets remained proportional across feature bins 5-50, regardless of feature dimensionality (**Figure 8G-J**). Slightly less agreement was observed in the top feature selection for datasets with continuous responses when compared to binary, where the latter demonstrated higher and more proportionate overlap (**Figure 8G-J**).

To further investigate how different data types influence model behavior for concatenated data, we evaluated the proportion of taxa versus metabolite features selected across bins sizes 1-50 (**Supplemental Figure 3A**). For Elastic Net, taxa consistently dominated feature selection across all datasets, with top features in every bin originating from the taxa data, regardless of dimensionality or response type (**Supplemental Figure 3B-E**). This suggests Elastic Net may be more sensitive to high dimensional features within concatenated models. In contrast, Random Forest and XGBoost selected a more balanced mix of metabolites and taxa features across bins, indicating a more even consideration of both data modalities in feature importance rankings (**Supplemental Figure 3B-E**).

Discussion

This study presents a comprehensive evaluation of machine learning performance and feature selection stability across multi-omics microbiome and metabolomics datasets using Elastic Net, Random Forest and XGBoost. By systematically comparing these models across both continuous and binary response variables, full dimensional and feature reduced datasets, and integration strategies including concatenation and ensemble-based methods, we aimed to uncover generalizable trends in model behavior and feature selection stability (**Figure 1**). One of the most consistent findings was the robust performance of Random Forest, particularly when combined with Non-Negative Least Squares (NNLS) integration, for datasets with continuous outcomes (**Figure 2**). This is consistent with prior studies highlighting the capacity of Random Forest to handle noisy, sparse, nonlinear and high dimensional data, such as microbiome data, effectively[5,27,28]. NNLS further enhanced model performance, reinforcing recent evidence that integration methods leveraging weighted combinations of predictions can outperform simpler or more complex stacking strategies for microbiome multi-omics data[16].

Across binary outcomes, Random Forest also outperformed other models, though the advantage was less pronounced with more comparable performances across the stacked integration methods (**Figure 3**). Notably, several individual 'omics models outperformed integrative models for both binary and continuous

responses, particularly metabolomics models. This suggests that in certain cases, single 'omics models may

identify strong predictive features and/or relationships that are lost during integration. We also observed that

model performance results were largely consistent between full dimensional and feature reduced datasets, with

Random Forest maintaining top performance across both response types (**Figure 2-3**). This suggests that,

despite substantial reductions in input features, overall predictive performance can be preserved when robust

models and integration strategies are applied[29].

In terms of feature selection, a key trend that emerged across all datasets is that Elastic Net consistently

identified a distinct and largely non-overlapping set of features compared to Random Forest and XGBoost

(**Figure 4-6**). This divergence is likely attributable to the linear and sparsity-inducing nature of Elastic Net, which

tends to select a small subset of features that are linearly associated with the outcome and includes groups of

correlated predictors. Alternatively, tree-based models are more adept at capturing nonlinear, higher-order

interactions and distributed signals[30].

This trend was particularly evident in the metabolomics datasets, where Random Forest and XGBoost

frequently shared more than 50% of their top-ranked features across bins, while Elastic Net consistently selected

a largely distinct set of features (**Figure 4**). The divergence in feature selection became even more pronounced

in the concatenated models that combined metabolomics and taxonomic data (**Figures 5–6**). Feature selection

stability was also more sensitive to dimensionality, with Elastic Net showing minimal overlap in top-ranked

features between full and feature-reduced datasets, especially in metabolomics, whereas Random Forest and

XGBoost demonstrated moderate consistency (**Figure 7**). These findings suggest that dimensionality reduction

alters the feature landscape more substantially for linear models, and that feature interpretability and

reproducibility in Elastic Net may be more influenced by preprocessing choices compared to tree-based

approaches.

When comparing feature selection between single 'omics and concatenated models, tree-based methods

retained a relatively high degree of consistency, while Elastic Net rarely recovered the same features—especially

in metabolomics analyses (**Figure 8**). Interestingly, for taxonomic features, Elastic Net showed much higher

agreement between concatenated and single 'omics models. This pattern likely reflects higher dimensionality in

the concatenated datasets, where taxa features outnumber metabolite features. As a result, Elastic Net appears

biased towards selecting taxa features, while Random Forest and XGBoost maintained a more balanced selection across both data types. These results are consistent with previous findings suggesting that tree-based models are more resilient to greater numbers of features and sparsity in high-dimensional data[30].

While this work provides several new insights into model behavior and feature selection for microbiome multi-omics data, there are several limitations. First, the study was conducted using four multi-omics microbiome datasets[11,14,25,26], limiting generalizability. Broader applications across a wider range of microbiome-related diseases would help validate these findings. Additionally, taxonomic profiling was performed using Kracken v2.1.1 and Bracken v2.8, which may produce different results when compared to other classification tools[31,32]. We also only selected one method for feature reduction analyses. We used Limma[33] for metabolomics data and Wilcoxon signed-rank test for the taxonomic relative abundance data. These might not be the best method for feature reduction and/or may not accurately reflect feature signal structure across all datasets. Exploration into other methods of microbiome data transformation and its effect on multi-omics data integration should be explored[34]. Furthermore, the concatenated datasets suffered from greater dimensionality, with taxa feature outnumbering metabolomics for all datasets. This imbalance likely contributed to Elastic Net's consistent selection of taxa-only features. Finally, a uniform hyperparameter grid was used across all models, which provided standardization but may not have captured optimal tuning for individual datasets. Future works should incorporate additional datasets exploring diverse diseases, consider additional integration methods such as DIABLO[35] or MOFA+[36], explore alternative dimensionality reduction strategies and taxa classification strategies and include refined hyperparameter optimization. Additionally, biological validation of selected features, through literature mining or experimental follow-up, is crucial to ensure biological relevance. Overall, this study offers a robust framework for evaluating machine learning performance and feature selection in multi-omics microbiome data, highlighting the strengths and limitations of current modeling strategies and laying the groundwork for more interpretable and scalable integrative analysis.

Key Points

- A total of 1,323 models were developed to comprehensively evaluate prediction performance and the robustness of feature selection using metabolomics and taxonomy datasets. These models included three widely used machine learning algorithms – Elastic Net, Random Forest and XGBoost – applied

across five integration strategies and single 'omics approaches on datasets with binary and continuous outcomes.

- For continuous outcomes, Random Forest combined with NNLS integration achieved the highest performance and maintained strong predictive performance across full dimensional and feature reduced datasets.

- For binary outcomes, Random Forest consistently performed well regardless of the integration strategy. Notably, single 'omics models, especially those using metabolomics data, outperformed integrative approaches.

- Tree-based models demonstrated greater consistency in feature selection across different dimensionalities and integration strategies.

Data Availability

Data is free and available through: https://github.com/borenstein-lab/microbiome-metabolome-curated-data/wiki/Data-overview#datasets-included[33]. Code for this integrative analysis and associated analysis are free and available through: https://github.com/suziepalmer10/Multiomics-Integrative-Pipeline.

Funding

This work is supported by National Institute of Allergy and Infectious Disease [grant numbers AI179406 to A.Y.K., AI169298 to X.Z., 5T32AI005284-43 to S.N.P.]; the University of Texas Southwestern Medical Center and Children's Health Cellular and ImmunoTherapeutics Program (to A.Y.K.); National Human Genome Research Institute [grant number HG011035 to D.L.]; and National Institute of General Medical Sciences [grant number GM126479 to X.Z.].

Competing Interests

AYK is a consultant for Prolacta Bioscience. AYK receives research funding from Novartis. AYK is co-founder of Aumenta Biosciences.

Table 1: Details for Data With <u>Continuous</u> Response Variables

Study Name (Subset)	Response Variable	Sample Size	Featured Reduced Details		Run Time (HH:MM)			Full Dimensional Details		Run Time (HH:MM)		
			Number of Metabolites	Number of Taxa	Elastic Net	Random Forest	XGBoost	Number of Metabolites	Number of Taxa	Elastic Net	Random Forest	XGBoost
Erawijantari	Age	96	71	775	00:09	00:01	01:55	524	3103	00:25	00:05	03:16
Erawijantari	Total Cholesterol	82	71	775	00:09	00:01	01:50	524	3103	00:25	00:06	02:48
Erawijantari	Glucose	84	71	775	00:09	00:01	01:47	524	3103	00:25	00:06	02:40
Franzosa (Control vs. CD)	Age	144	2175	3603	00:30	00:13	04:50	8848	5137	01:34	00:47	09:56
Franzosa (Control vs. CD)	Fecal Calprotectin (Fp)	107	2175	3603	00:30	00:10	04:10	8848	5137	01:06	00:32	06:50
Franzosa (Control vs. IBD)	Age	220	2519	3849	00:39	00:24	06:15	8848	5137	01:18	01:00	10:11
Franzosa (Control vs IBD)	Fecal Calprotectin (Fp)	153	2519	3849	00:33	00:16	04:54	8848	5137	01:08	00:42	07:48
Franzosa (Control vs. UC)	Age	132	1268	1816	00:18	00:06	03:43	8848	5137	01:31	00:45	09:25
Franzosa (Control vs. UC)	Fecal Calprotectin (Fp)	88	1268	1816	00:16	00:04	03:01	8848	5137	01:25	00:36	06:33
Wang	Age	287	71	2261	00:13	00:09	03:22	276	5293	00:26	00:26	04:50
Wang	BMI	280	71	2261	–	00:09	03:06	276	5293	00:26	00:27	04:26
Wang	Creatinine	281	71	2261	00:12	00:10	02:49	276	5293	00:25	00:30	04:10
Wang	eGFR	281	71	2261	00:12	00:11	03:37	276	5293	00:25	00:32	05:11
Wang	Urea	281	71	2261	00:13	00:09	03:33	276	5293	00:25	00:27	04:59
Yachida (All Samples)	Age	347	173	1823	00:23	00:11	03:22	450	5770	00:32	00:34	05:27
Yachida (All Samples)	Alcohol	347	173	1823	00:24	00:11	03:43	450	5770	00:31	00:33	05:47
Yachida (All Samples)	BMI	347	173	1823	–	00:11	03:05	450	5770	–	00:35	05:05
Yachida (All Samples)	Brinkman Index	347	173	1823	00:25	00:12	03:41	450	5770	00:32	00:36	05:46

'-' represents models that failed to complete.

Table 2: Details for Data With <u>Binary</u> Response Variables

Study Name (Subset)	Response Variable	Sample Size	Featured Reduced Details					Full Dimensional Details				
			Number of Metabolites	Number of Taxa	Run Time (HH:MM) Elastic Net	Random Forest	XGBoost	Number of Metabolites	Number of Taxa	Run Time (HH:MM) Elastic Net	Random Forest	XGBoost
Erawijantari	No Alcohol vs. Alcohol	96 (47/49)	71	775	00:09	00:01	01:55	524	3103	00:25	00:05	03:16
Erawijantari	Female vs. Male	96 (32/64)	71	775	–	00:04	01:16	524	3103	00:31	00:28	02:14
Erawijantari	Control vs. Gastrectomy	96 (54/42)	71	775	00:11	00:04	01:08	524	3103	00:23	00:27	02:09
Franzosa (Control vs. CD)	Control vs. CD	144 (56/88)	2175	3603	00:38	01:13	02:43	8848	5137	01:52	04:52	05:29
Franzosa (Control vs. IBD)	Control vs. IBD	220 (56/164)	2519	3849	00:54	02:19	03:16	8848	5137	01:48	07:34	05:34
Franzosa (Control vs. UC)	Control vs. UC	132 (56/76)	1268	1816	00:22	00:30	02:07	8848	5137	01:53	04:45	05:21
Wang	Female vs. Male	287 (107/180)	71	2261	–	01:03	02:09	276	5293	–	03:49	03:06
Wang	Control vs. Kidney Failure	287 (67/220)	71	2261	00:17	00:50	01:36	276	5293	00:34	03:11	02:29
Yachida (All Samples)	Female vs. Male	347 (143/204)	173	1823	00:18	01:08	02:12	450	5770	00:42	05:16	03:39
Yachida (Control vs MP & CRC)	Control vs. MP and CRC	277 (127/150)	173	1823	00:16	00:51	01:58	450	5770	00:35	04:02	03:15
Yachida (Control vs MP & Stage 0)	Control vs. MP & Stage 0 CRC	194 (127/67)	173	1823	00:07	00:09	01:24	450	5770	–	02:28	02:46
Yachida (Control vs Stage 1 & 2)	Control vs. Stage 1 & 2 CRC	196 (127/67)	173	1823	00:05	00:05	01:13	450	5770	–	02:30	02:47
Yachida (Control vs Stage 3 & 4)	Control vs. Stage 3 & 4 CRC	181 (127/54)	173	1823	00:07	00:08	01:15	450	5770	–	02:08	02:34

'-' represents models that failed to complete.

A) Data Processing

Taxonomy

Taxa Minimum Threshold Filtering

Differential Analysis to Reduce Features

Metabolomics

Scaling Based on Data Distribution

Differential Analysis to Reduce Features

Response

Remove Missing Values

Full Dimensional Datasets

Feature Reduced Datasets

B) Model Training and Testing

Training (80%)

fold 1 ••• fold *n*

v-fold cross validation

Performance

Input Models: Elastic Net, Random Forest, XGBoost
Integration Techniques: Single Omics (Taxonomy and Metabolomics), Concatenatation, Averaged Stacking, Weighted Non-Negative Least Squares (NNLS), Lasso Stacking, Partial Least Squares.

C) Feature Evaluation

Feature Comparisons

Feature Importance

1.
2.
⋮
n

Comparisons:
Elastic Net vs. Random Forest vs. XGBoost
Full Dimensionall vs. Feature Reduced Datasets
Concatenated vs. Taxonomy/Metabolomics

 Figure Legends

Figure 1. Pipeline for the Multi-omics Integrative Analysis. A) Matched taxonomy and metabolomics data

are processed separately. Taxa with an abundance below 0.001 across all samples are removed. Depending on

the study design and the method used for metabolomics, data are appropriately scaled. Response variables with

missing values are excluded from downstream analyses. Differential expression analyses are conducted

separately for metabolomics (indicated in red) and taxa (indicated in blue). Both the original datasets (taxa – 1,

metabolomics – 1) and their corresponding response variables are concatenated. Similarly, the feature reduced

datasets (taxa – 2, metabolomics – 2) and their response variables are concatenated, enabling comparisons

between original and feature-reduced data. **B)** The processed datasets are then analyzed using the multi-omics

integration pipeline. The data are first split into 80% training and 20% testing sets. The training data undergo 3

× 5-fold cross-validation to optimize model performance. Performance metrics—area under the receiver

operating characteristic (ROC) curve (AUROC) for binary outcomes and root mean squared error (RMSE) for

continuous outcomes—are used to evaluate both validation and testing performance. The base machine learning

models are Elastic Net, XGBoost, and Random Forest. Integration strategies include simple concatenation,

averaged stacking, weighted non-negative least squares (NNLS), Lasso stacking, and Partial Least Squares

(PLS). Additionally, single 'omics models are generated separately for taxonomy and metabolomics datasets.

For each response variable, 21 models are generated (or attempted), representing various combinations of base

learners, integration methods, and data types. **C)** After model generation, feature extraction is performed. Top

features selected by each machine learning model (Elastic Net, Random Forest, XGBoost) are compared to

identify consistently important variables. Further comparisons are conducted to assess feature selection

differences between full dimensional and feature reduced datasets. Lastly, the features selected from

concatenated taxonomy and metabolomics base models are compared to those identified in the integrated

concatenated models to evaluate the impact of data integration on feature importance.

A) Performance of Feature Reduced Data with Continuous Response Variables

B) Feature Reduced Categorical Distributions with the Lowest RMSE

Top Model
- Elastic Net (24.1%)
- Random Forest (44.1%)
- XGBoost (31.5%)

Top Integration Method
- Concatenated (3.7%)
- Averaged Stacked (AS) (5.6%)
- Weighted NNLS (W-NNLS) (72.2%)
- Lasso Stacked (LS) (11.1%)
- PLS (7.4%)

C) Performance of Full Dimensional Data with Continuous Response Variables

D) Full Dimensional Categorical Distributions with the Lowest RMSE

Top Model
- Elastic Net (24.1%)
- Random Forest (46.3%)
- XGBoost (29.6%)

Top Integration Method
- Concatenated (0.0%)
- Averaged Stacked (AS) (7.4%)
- Weighted NNLS (W-NNLS) (75.9%)
- Lasso Stacked (LS) (9.3%)
- PLS (7.4%)

Figure 2. Testing Performance Metrics Across Continuous Feature Reduced and Full Dimensional

Datasets. Testing performance for feature reduced matched microbiome and metabolomics datasets with

continuous response variables are represented by **A)**. The top performing machine learning models and top

integration methods for each of the feature reduced datasets are represented as pie charts in **B)**. Testing

performance for full dimensional matched microbiome and metabolomics datasets with continuous response

variables are represented by **C)**. The top performing machine learning models and top integration methods for

each of the full dimensional datasets are represented as pie charts in **D)**. Performance is assessed using RMSE

Mean. Actual RMSE values are shown in each cell of the heatmap with $\log_{10}$ scaling for the coloring. Response

variables used for each model are shown on the right of each heatmap. The top performing models per response

variable are represented by '*', where '*' is the best performing model, '**' is the second-best performing model

and '***' is the third-best performing model. The proportions are calculated using the top three best performing

models (*, **, ***) per response variable. 'T' represents ties in performance; response variables with majority 'T'

were not used for pie chart calculations. Individual 'omics (Metabolomics and MSS) were not represented in the

pie charts. '>' or 'T' represents increased or tied performance of an individual 'omics to one of the top performing

integration models ('*', '**', and '***').

A) Performance of Feature Reduced Data with Continuous Response Variables

C) Performance of Full Dimensional Data with Continuous Response Variables

B) Feature Reduced Categorical Distributions with the Lowest RMSE

D) Full Dimensional Categorical Distributions with the Lowest RMSE

Supplemental Figure 1. **Validation Performance Metrics Across Continuous Feature Reduced and Full Dimensional Datasets.** Validation performance for feature reduced matched microbiome and metabolomics datasets with continuous response variables are represented by **A)**. The top performing machine learning models and top integration methods for each of the feature reduced datasets are represented as pie charts in **B)**. Validation performance for full dimensional matched microbiome and metabolomics datasets with continuous response variables are represented by **C)**. The top performing machine learning models and top integration methods for each of the full dimensional datasets are represented as pie charts in **D)**. Performance calculated using RMSE Mean. Actual RMSE values are shown in each cell of the heatmap with $\log_{10}$ scaling for the coloring. Response variables used for each model are shown on the right of each heatmap. The top performing models per response variable are represented by '*', where '*' is the best performing model, '**' is the second-best performing model and '***' is the third-best performing model. The proportions are calculated using the top three best performing models (*, **, ***) per response variable. 'T' represents ties in performance; response variables with majority 'T' were not used for pie chart calculations. Individual 'omics (Metabolomics and MSS) were not represented in the pie charts. '>' or 'T' represents increased or tied performance of an individual 'omics to one of the top performing integration models ('*', '**', and '***').

A) Performance of Feature Reduced Data with Binary Response Variables

B) Feature Reduced Categorical Distributions with the Lowest AUROC

C) Performance of Full Dimensional Data with Binary Response Variables

D) Full Dimensional Categorical Distributions with the Lowest AUROC

Figure 3. Test Performance Metrics Across Binary Feature Reduced and Full Dimensional Datasets.
Testing performance for feature reduced matched microbiome and metabolomics datasets with binary response variables are represented by **A)**. The top performing machine learning models and top integration methods for each of the feature reduced datasets are represented as pie charts in **B)**. Testing performance for full dimensional matched microbiome and metabolomics datasets with binary response variables are represented by **C)**. The top performing machine learning models and top integration methods for each of the full dimensional datasets are represented as pie charts in **D)**. Performance is calculated using AUROC mean, which are shown in the heatmap. Response variables used for each model are shown on the right of each heatmap. The top performing models per response variable are represented by '*', where '*' is the best performing model, '**' is the second-best performing model and '***' is the third-best performing model. The proportions are calculated using the top three best performing models (*, **, ***) per response variable. 'T' represents ties in performance; response variables with majority 'T' were not used for pie chart calculations. Individual 'omics (Metabolomics and MSS) were not represented in the pie charts. '>' or 'T' represents increased or tied performance of an individual 'omics to one of the top performing integration models ('*', '**', and '***').

A) Performance of Feature Reduced Data with Binary Response Variables

C) Performance of Full Dimensional Data with Binary Response Variables

B) Feature Reduced Categorical Distributions with the Lowest AUROC

Top Model
- Elastic Net (19.4%)
- Random Forest (55.6%)
- XGBoost (25.0%)

Top Integration Method
- Concatenated (0.0%)
- Averaged Stacked (AS) (11.1%)
- Weighted NNLS (W-NNLS) (25.0%)
- Lasso Stacked (LS) (27.8%)
- PLS (36.1%)

D) Full Dimensional Categorical Distributions with the Lowest AUROC

Top Model
- Elastic Net (38.9%)
- Random Forest (19.4%)
- XGBoost (41.7%)

Top Integration Method
- Concatenated (2.8%)
- Averaged Stacked (AS) (8.3%)
- Weighted NNLS (W-NNLS) (25.0%)
- Lasso Stacked (LS) (27.8%)
- PLS (36.1%)

Supplemental Figure 2. Validation Performance Metrics Across Binary Feature Reduced and Full Dimensional Datasets. Validation performance for feature reduced matched microbiome and metabolomics datasets with binary response variables are represented by **A)**. The top performing machine learning models and top integration methods for each of the feature reduced datasets are represented as pie charts in **B)**. Testing performance for full dimensional matched microbiome and metabolomics datasets with binary response variables are represented by **C)**. The top performing machine learning models and top integration methods for each of the full dimensional datasets are represented as pie charts in **D)**. Performance is calculated using AUROC mean, which are shown in the heatmap. Response variables used for each model are shown on the right of each heatmap. The top performing models per response variable are represented by '*', where '*' is the best performing model, '**' is the second-best performing model and '***' is the third-best performing model. The proportions are calculated using the top three best performing models (*, **, ***) per response variable. 'T' represents ties in performance; response variables with majority 'T' were not used for pie chart calculations. Individual 'omics (Metabolomics and MSS) were not represented in the pie charts. '>' or 'T' represents increased or tied performance of an individual 'omics to one of the top performing integration models ('*', '**', and '***').

A) Metabolite Comparisons for Feature Reduced and Full Dimensional Binary and Continuous Datasets

Extract Elastic Net (**ENET**), Random Forest (**RF**) and XGBoost (**XGB**) Base Model Features

Feature Normalization and Scaling

Absolute Value
Min-Max Normalization

Feature Comparison Between Machine Learning Models

Violin Plot Legend

○ Proportion of Overlapping or Unique Metabolites Selected by ENET, RF and/or XGB for Each Dataset

B) Comparisons of Metabolites Across Feature Reduced Datasets With Continuous Response Variables

C) Comparisons of Metabolites Across Full Dimensional Datasets With Continuous Response Variables

D) Comparisons of Metabolites Across Feature Reduced Datasets With Binary Response Variables

E) Comparisons of Metabolites Across Full Dimensional Datasets With Binary Response Variables

Figure 4. Metabolite Feature Selection Comparisons across Elastic Net, Random Forest and XGBoost. The pipeline for extracting feature importance metrics and performing comparisons are represented in **A)**. Feature Importance values for Elastic Net (ENET), Random Forest (RF) and XGBoost (XGB) were extracted for each cross-validation step. The absolute value was taken for each value and min-max normalization was performed. The overall average for each feature was then used as the metric for feature importance. See **Methods** for more details. Comparisons for the metabolites selected across feature reduced datasets with continuous response variables are represented by **B)**. Comparisons for the metabolites selected across full dimensional datasets with continuous response variables are represented by **C)**. Comparisons for the metabolites selected across feature reduced datasets with binary response variables are represented by **D)**. Comparisons for the metabolites selected across full dimensional datasets with binary response variables are represented by **E)**. For each analysis, top features for ENET, RF and XGB were binned into 5 levels – top feature (i.e., features with highest feature importance values), top 5 features, top 10 feature, top 20 features and top 50 features. For each bin, comparisons for the features selected between ENET, RF and XGB were performed to identify joint and unique features. The distribution of dataset proportions for overlapping and unique features are represented by violin plots for panels **B-E**.

A) Taxa Comparisons for Feature Reduced and Full Dimensional Binary and Continuous Datasets

Extract Elastic Net (**ENET**), Random Forest (**RF**) and XGBoost (**XGB**) Base Model Features

Feature Normalization and Scaling

Absolute Value
Min-Max Normalization

Feature Comparison Between Machine Learning Models

Violin Plot Legend

Proportion of Overlapping or Unique Taxa Selected by ENET, RF and/or XGB for Each Dataset

B) Comparisons of Taxa Across Feature Reduced Datasets With Continuous Response Variables

Overlap ENET, RF and XGB | Overlap ENET and RF | Overlap ENET and XGB | Overlap RF and XGB | Unique to ENET | Unique to RF | Unique to XGB

Top Feature
Top 5
Top 10
Top 20
Top 50

Proportion of Overlapping Features

C) Comparisons of Taxa Across Full Dimensional Datasets With Continuous Response Variables

Overlap ENET, RF and XGB | Overlap ENET and RF | Overlap ENET and XGB | Overlap RF and XGB | Unique to ENET | Unique to RF | Unique to XGB

Top Feature
Top 5
Top 10
Top 20
Top 50

Proportion of Overlapping Features

D) Comparisons of Taxa Across Feature Reduced Datasets With Binary Response Variables

Overlap ENET, RF and XGB | Overlap ENET and RF | Overlap ENET and XGB | Overlap RF and XGB | Unique to ENET | Unique to RF | Unique to XGB

Top Feature
Top 5
Top 10
Top 20
Top 50

Proportion of Overlapping Features

E) Comparisons of Taxa Across Full Dimensional Datasets With Binary Response Variables

Overlap ENET, RF and XGB | Overlap ENET and RF | Overlap ENET and XGB | Overlap RF and XGB | Unique to ENET | Unique to RF | Unique to XGB

Top Feature
Top 5
Top 10
Top 20
Top 50

Proportion of Overlapping Features

Figure 5. Taxa Feature Selection Comparisons across Elastic Net, Random Forest and XGBoost. The pipeline for extracting feature importance metrics and performing comparisons are represented in **A)**. Feature Importance values for Elastic Net (ENET), Random Forest (RF) and XGBoost (XGB) were extracted for each cross-validation step. The absolute value was taken for each value and min-max normalization was performed. The overall average for each feature was then used as the metric for feature importance. See **Methods** for more details. Comparisons for the taxa selected across feature reduced datasets with continuous response variables are represented by **B)**. Comparisons for the taxa selected across full dimensional datasets with continuous response variables are represented by **C)**. Comparisons for the taxa selected across feature reduced datasets with binary response variables are represented by **D)**. Comparisons for the taxa selected across full dimensional datasets with binary response variables are represented by **E)**. For each analysis, top features for ENET, RF and XGB were binned into 5 levels – top feature (i.e., features with highest feature importance values), top 5 features, top 10 feature, top 20 features and top 50 features. For each bin, comparisons for the features selected between ENET, RF and XGB were performed to identify joint and unique features. The distribution of dataset proportions for overlapping and unique features are represented by violin plots for panels **B-E**.

A) Concatenation (Taxa and Metabolite) Comparisons for Feature Reduced and Full Dimensional Binary and Continuous Datasets

Extract Elastic Net **(ENET)**, Random Forest **(RF)** and XGBoost **(XGB)** Base Model Features

Feature Comparison Between Machine Learning Models

Feature Normalization and Scaling

Absolute Value
Min-Max Normalization

Violin Plot Legend

Proportion of Overlapping or Unique Taxa and Metabolites Selected by ENET, RF and/or XGB for Each Dataset

B) Comparisons of Taxa and Metabolites Across Feature Reduced Datasets With Continuous Response Variables

C) Comparisons of Taxa and Metabolites Across Full Dimensional Datasets With Continuous Response Variables

D) Comparisons of Taxa and Metabolites Across Feature Reduced Datasets With Binary Response Variables

E) Comparisons of Taxa and Metabolites Across Full Dimensional Datasets With Binary Response Variables

Figure 6. Concatenation Feature Selection Comparisons across Elastic Net, Random Forest and
XGBoost. The pipeline for extracting feature importance metrics and performing comparisons are represented
in **A)**. Feature Importance values for Elastic Net (ENET), Random Forest (RF) and XGBoost (XGB) were
extracted for each cross-validation step. The absolute value was taken for each value and min-max normalization
was performed. The overall average for each feature was then used as the metric for feature importance. See
Methods for more details. Comparisons for the metabolites and taxa selected across feature reduced datasets
with continuous response variables are represented by **B)**. Comparisons for the metabolites and taxa selected
across full dimensional datasets with continuous response variables are represented by **C)**. Comparisons for the
metabolites and taxa selected across feature reduced datasets with binary response variables are represented
by **D)**. Comparisons for the metabolites and taxa selected across full dimensional datasets with binary response
variables are represented by **E)**. For each analysis, top features for ENET, RF and XGB were binned into 5 levels
– top feature (i.e., features with highest feature importance values), top 5 features, top 10 feature, top 20 features
and top 50 features. For each bin, comparisons for the features selected between ENET, RF and XGB were
performed to identify joint and unique features. The distribution of dataset proportions for overlapping and unique
features are represented by violin plots for panels **B-E**.

A) Feature Selection Comparisons Between Feature Reduced (FR) and Full Dimensional (FD) Datasets

Extract Elastic Net **(ENET)**, Random Forest **(RF)** and XGBoost **(XGB)** Base Model Features

Feature Normalization and Scaling

Absolute Value
Min-Max Normalization

Feature Comparison Between Machine Learning Models

FR   FD

Violin Plot Legend
○ ENET
○ RF
○ XGB

○ Proportion of Overlapping or Unique Taxa and Metabolites Selected by ENET, RF or XGB for Each Dataset

B) Comparisons of Metabolite Features Selected Between FR and FD Datasets with Continuous Response Variables

C) Comparisons of Taxa Features Selected Between FR and FD Datasets with Continuous Response Variables

D) Comparisons of Concatenated Features Selected Between FR and FD Datasets with Continuous Response Variables

E) Comparisons of Metabolite Features Selected Between FR and FD Datasets with Binary Response Variables

F) Comparisons of Taxa Features Selected Between FR and FD Datasets with Binary Response Variables

G) Comparisons of Concatenated Features Selected Between FR and FD Datasets with Binary Response Variables

Figure 7. Feature Selection Comparisons Between Featured Reduced (FR) and Full Dimensional (FD)

Datasets Across Elastic Net (ENET), Random Forest (RF) and XGBoost (XGB) Models. The pipeline for

extracting feature importance metrics and performing comparisons are represented in **A)**. Feature Importance

values for ENET, RF and XGB were extracted for each cross-validation step. The absolute value was taken for

each value and min-max normalization was performed. The overall average for each feature was then used as

the metric for feature importance. See **Methods** for more details. Comparisons for the metabolites selected by

feature reduced and full dimensional datasets with continuous response variables are represented by **B)**.

Comparisons for the taxa selected by feature reduced and full dimensional datasets with continuous response

variables are represented by **C)**. Comparisons for the taxa and metabolites selected by feature reduced and full

dimensional datasets with continuous response variables are represented by **D)**. Comparisons for the

metabolites selected by feature reduced and full dimensional datasets with binary response variables are

represented by **E)**. Comparisons for the taxa selected by feature reduced and full dimensional datasets with

binary response variables are represented by **F)**. Comparisons for the taxa and metabolites selected by feature

reduced and full dimensional datasets with continuous response variables are represented by **G)**. For each

analysis, top features for ENET, RF and XGB were binned into 5 levels – top feature (i.e., features with highest

feature importance values), top 5 features, top 10 feature, top 20 features and top 50 features. For each bin,

comparisons for the features selected by feature reduced and full dimensional datasets for ENET, RF and XGB

were performed to identify joint and unique features. The distribution of dataset proportions for overlapping and

unique features are represented by violin plots for panels **B-G**.

A) Feature Selection Comparisons Between Concatenated and Metabolite Models

B) Comparisons of Overlapping Metabolite Features Selected by Concatenated and Metabolite Models for Feature Reduced Datasets with Continuous Response Variables

C) Comparisons of Overlapping Metabolite Features Selected by Concatenated and Metabolite Models for Full Dimensional Datasets with Continuous Response Variables

D) Comparisons of Overlapping Metabolite Features Selected by Concatenated and Metabolite Models for Feature Reduced Datasets with Binary Response Variables

E) Comparisons of Overlapping Metabolite Features Selected by Concatenated and Metabolite Models for Full Dimensional Datasets with Binary Response Variables

F) Feature Selection Comparisons Between Concatenated and Taxa Models

G) Comparisons of Overlapping Taxa Features Selected by Concatenated and Taxa Models for Feature Reduced Datasets with Continuous Response Variables

H) Comparisons of Overlapping Taxa Features Selected by Concatenated and Taxa Models for Full Dimensional Datasets with Continuous Response Variables

I) Comparisons of Overlapping Taxa Features Selected by Concatenated and Taxa Models for Feature Reduced Datasets with Binary Response Variables

J) Comparisons of Overlapping Taxa Features Selected by Concatenated and Taxa Models for Full Dimensional Datasets with Binary Response Variables

Figure 8. Feature Selection Comparisons Between Concatenated (Concat) and Individual (Metabolomics/Taxa) Elastic Net (ENET), Random Forest (RF) and XGBoost (XGB) Models. The pipeline for extracting feature importance metrics and performing comparisons between concatenated and metabolites models are represented in **A)**. Feature Importance values for ENET, RF and XGB were extracted for each cross-validation step. The absolute value was taken for each value and min-max normalization was performed. The overall average for each feature was then used as the metric for feature importance. See **Methods** for more details. Comparisons for the metabolites selected by concatenated and metabolomics models for feature reduced datasets with continuous response variables are represented by **B)**. Comparisons for the metabolites selected by concatenated and metabolomics models for full dimensional datasets with continuous response variables are represented by **C)**. Comparisons for the metabolites selected by concatenated and metabolomics models for feature reduced datasets with binary response variables are represented by **D)**. Comparisons for the metabolites selected by concatenated and metabolomics models for full dimensional datasets with binary response variables are represented by **E)**. The pipeline for extracting feature importance metrics and performing comparisons between concatenated and metabolites models are represented in **F)** (See details above (**A**) and in **Methods**). Comparisons for the taxa selected by concatenated and taxa models for feature reduced datasets with continuous response variables are represented by **G)**. Comparisons for the taxa selected by concatenated and taxa models for full dimensional datasets with continuous response variables are represented by **H)**. Comparisons for the taxa selected by concatenated and taxa models for feature reduced datasets with binary response variables are represented by **I)**. Comparisons for the taxa selected by concatenated and taxa models for full dimensional datasets with binary response variables are represented by **J)**. For each analysis, top features for ENET, RF and XGB were binned into 5 levels – top feature (i.e., features with highest feature importance values), top 5 features, top 10 feature, top 20 features and top 50 features. For each bin, comparisons for the features selected by concatenated and individual (metabolomics and taxa) 'omics models for ENET, RF and XGB were performed to identify joint and unique features. The distribution of dataset proportions for overlapping and unique features are represented by violin plots for panels **B-E** and **G-J**.

A) Proportions of Taxa and Metabolite Features Selected by Concatenated Models

B) Proportions of Metabolites and Taxa Selected as Top Features by Feature Reduced Concatenated Models Using Continuous Response Variables

C) Proportions of Metabolites and Taxa Selected as Top Features by Full Dimensional Concatenated Models Using Continuous Response Variables

D) Proportions of Metabolites and Taxa Selected as Top Features by Feature Reduced Concatenated Models Using Binary Response Variables

E) Proportions of Metabolites and Taxa Selected as Top Features by Full Dimensional Concatenated Models Using Binary Response Variables

Supplemental Figure 3: Proportions of Taxa and Metabolite Features Selected by Concatenated Models.
The pipeline for extracting feature importance metrics and calculating proportions of the top features that are metabolites and taxa selected by the concatenated models are represented by **A)**. The Proportions of taxa and metabolite features selected by ENET, RF and XGB for concatenated feature reduced datasets with continuous response variable are shown in **B)**. The Proportions of taxa and metabolite features selected by ENET, RF and XGB for concatenated full dimensional datasets with continuous response variable are shown in **C)**. The Proportions of taxa and metabolite features selected by ENET, RF and XGB for concatenated feature reduced datasets with binary response variable are shown in **D)**. The Proportions of taxa and metabolite features selected by ENET, RF and XGB for concatenated full dimensional datasets with binary response variable are shown in **E)**. The distribution of proportions for each dataset are represented as violin plots. For each analysis, top features for ENET, RF and XGB were binned into 5 levels – top feature (i.e., features with highest feature importance values), top 5 features, top 10 feature, top 20 features and top 50 features. The red violin plot represents metabolomics proportions for each dataset and the blue violin plot represents the taxa proportions for each dataset. These violin plot values are inverse proportional for each dataset, so that if the proportion of metabolites for Dataset A = 1, then the proportion of taxa for Dataset A = 0.

References

1. Gilbert JA, Quinn RA, Debelius J, et al. Microbiome-wide association studies link dynamic microbial consortia to disease. Nature. 2016 Jul 7;535(7610).

2. Joos R, Boucher K, Lavelle A, et al. Examining the healthy human microbiome concept. Nature Reviews Microbiology. 2024 Mar;23(3):192-205.

3. Walter J, Armet AM, Finlay BB, et al. Establishing or Exaggerating Causality for the Gut Microbiome: Lessons from Human Microbiota-Associated Rodents. Cell. 2020 Jan 23;180(2):221-232.

4. Corander J, Hanage WP, Pensar J. Causal discovery for the microbiome. The Lancet Microbe. 2022 Nov;3(11):e881-e887.

5. Knight R, Vrbanac A, Taylor BC, et al. Best practices for analysing microbiomes. Nature Reviews Microbiology. 2018 Jul;16(7):410-422.

6. Galloway-Peña J, Hanson B. Tools for Analysis of the Microbiome. Digestive Diseases and Sciences. 2020 Mar;65(3):674-685.

7. Song EJ, Lee ES, Nam YD. Progress of analytical tools and techniques for human gut microbiome research. Journal of Microbiology. 2018 Oct;56(10):693-705.

8. Hasin Y, Seldin M, Lusis A. Multi-omics approaches to disease. Genome Biology. 2017 May;83(18).

9. Chetty A, Blekhman R. Multi-omic approaches for host-microbiome data integration. Gut Microbes. 2024 Jan-Dec;16(1):2297860.

10. Li P, Luo H, Ji B, et al. Machine learning for data integration in human gut microbiome. Microbial Cell Factories. 2022 Nov 23;21(1):241.

11. Franzosa EA, Sirota-Madi A, Avila-Pacheco J, et al. Gut microbiome structure and metabolic activity in inflammatory bowel disease. Nature Microbiology. 2019 Feb;4(2):293-305.

12. Li J, Zhao F, Wang Y, et al. Gut microbiota dysbiosis contributes to the development of hypertension. Microbiome. 2017 Feb 1;5(1):14.

13. Gao B, Duan Y, Lang S, et al. Functional Microbiomics Reveals Alterations of the Gut Microbiome and Host Co-Metabolism in Patients With Alcoholic Hepatitis. Hepatology Communications. 2020 Jun 19;4(8):1168-1182.

14. Yachida S, Mizutani S, Shiroma H, et al. Metagenomic and metabolomic analyses reveal distinct stage-specific phenotypes of the gut microbiota in colorectal cancer. Nature Medicine. 2019 Jun;25(6):968-976.

15. Ghaemi MS, DiGiulio DB, Contrepois K, et al. Multiomics modeling of the immunome, transcriptome, microbiome, proteome and metabolome adaptations during human pregnancy. Bioinformatics. 2019 Jan 1;35(1):95-103.

16. Mallick H, Porwal A, Saha S, et al. An integrated Bayesian framework for multi-omics prediction and classification. Statistics in Medicine. 2024 Feb 28;43(5):983-1002.

17. Liu YX, Qin Y, Chen T, et al. A practical guide to amplicon and metagenomic analysis of microbiome data. Protein & Cell. 2021 May;12(5):315-330.

18. Poussin C, Sierro N, Boué S, et al. Interrogating the microbiome: experimental and computational considerations in support of study reproducibility. Drug Discovery Today. 2018 Sep;23(9):1644-1657.

19. Huttenhower C, Finn RK, McHardy AC. Challenges and opportunities in sharing microbiome data and analyses. Nature Microbiology. 2023 Nov;8(11):1960-1970.

20. Santiago-Rodriguez TM, Hollister EB. Multi 'omic data integration: A review of concepts, considerations, and approaches. Seminars in Perinatology. 2021 Oct;45(6):151456.

21. Papoutsoglou G, Tarazona S, Lopes MB, et al. Machine learning approaches in microbiome research: challenges and best practices. Frontiers in Microbiology. 2023 Sep 22;14:1261889.

22. Mullen KM, van Stokkum IHM. nnls: The Lawson-Hanson Algorithm for Non-Negative Least Squares (NNLS). https://rdrr.io/cran/nnls/ (27 June 2025, date last accessed).

23. Friedman J, Hastie T, Tibshirani R. et al. Regularization Paths for Generalized Linear Models via Coordinate Descent. Journal of Statistical Software. 2010 Aug;33(1):1-22.

24. Mevik B-H, Wehrens R. The pls Package: Principal Component and Partial Least Squares Regression in R. Journal of Statistical Software. 2007 Jan;18:1–23.

25. Erawijantari PP, Mizutani S, Shiroma H, et al. Influence of gastrectomy for gastric cancer treatment on faecal microbiome and metabolome profiles. Gut. 2020 Aug;69(8):1404-1415.

26. Wang X, Yang S, Li S, et al. Aberrant gut microbiota alters host metabolome and impacts renal failure in humans and rodents. Gut. 2020 Dec;69(12):2131-2142.

27. Hernández Medina R, Kutuzova S, Nielsen KN, et al. Machine learning and deep learning applications in microbiome research. ISME Communications. 2022 Oct 6;2(1):98.

28. Kumar B, Lorusso E, Fosso B, et al. A comprehensive overview of microbiome data in the light of machine learning applications: categorization, accessibility, and future directions. Frontiers in Microbiology. 2024 Feb 13;15:1343572.

29. Saeys Y, Inza I, Larrañaga P. A review of feature selection techniques in bioinformatics. Bioinformatics. 2007 Oct 1;23(19):2507-17.

30. Díaz-Uriarte R, Alvarez de Andrés S. Gene selection and classification of microarray data using random forest. BMC Bioinformatics. 2006 Jan 6;7:3.

31. Wood DE, Lu J, Langmead B. Improved metagenomic analysis with Kraken 2. Genome Biology. 2019 Nov 28;20(1):257.

32. Lu J, Breitwieser FP, Thielen P, et al. Bracken: estimating species abundance in metagenomics data. PeerJ Computer Science. 2017 Jan;3:e104.

33. Ritchie ME, Phipson B, Wu D, et al. Limma powers differential expression analyses for RNA-sequencing and microarray studies. Nucleic Acids Research. 2015 Apr 20;43(7):e47.

34. Karwowska Z, Aasmets O; Kosciolek T, et al. Effects of data transformation and model selection on feature importance in microbiome classification data. Microbiome. 2025 Jan 4;13(1):2.

35. Singh A, Shannon CP, Gautier B, et al. DIABLO: an integrative approach for identifying key molecular drivers from multi-omics assays. Bioinformatics. 2019 Sep 1;35(17):3055-3062.

36. Argelaguet R, Arnol D, Bredikhin D, et al. MOFA+: a statistical framework for comprehensive integration of multi-modal single-cell data. Genome Biology. 2020 May 11;21(1):111.

37. Muller E, Algavi YM, Borenstein E. The gut microbiome-metabolome dataset collection: a curated resource for integrative meta-analysis. Npj Biofilms and Microbiomes. 2022 Oct 15;8(1):79.

38. Wickham H, Averick M, Bryan J, et al. Welcome to the Tidyverse. Journal of Open Source Software. 2019 Nov; 4(43):1686.

39. Robin X, Turck N, Hainard A, et al. PROC: an open-source package for R and S+ to analyze and compare ROC curves. BMC Bioinformatics. 2011 Mar 17;12:77.

40. Gamer M, Lemon J, Fellows I, et al. irr: Various Coefficients of Interrater Reliability and Agreement. 2019. https://cran.r-project.org/web/packages/irr/index.html (27 June 2025, date last accessed).

41. Greener JG, Kandathil SM, Moffat L, et al. A guide to machine learning for biologists. Nature Reviews Molecular Cell Biology. 2022 Jan;23(1):40-55.

42. Li Y, Wu FX, Ngom A. A review on machine learning principles for multi-view biological data integration. Briefings in Bioinformatics. 2018 Mar 1;19(2):325-340.

43. Ghannam RB, Techtmann SM. et al. Machine learning applications in microbial ecology, human microbiome studies, and environmental monitoring. Computational and Structural Biotechnology Journal. 2021 Jan 27;19:1092-1107.

44. Kühn M. Building Predictive Models in R Using the caret Package. Journal of Statistical Software. 2008 Nov;28(5):1-26.

45. Liaw A, Wiener MC. Classification and Regression by randomForest. R News. 2002 Dec;2:18-22.

46. Chen T, Guestrin C. XGBOOST: A Scalable Tree Boosting System. KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. 2016 Aug 13; 785-794.

47. Gu Z. Complex heatmap visualization. iMeta. 2022 Sept;3(1):e43.

48. Gu Z, Eils R, Schlesner M. Complex heatmaps reveal patterns and correlations in multidimensional genomic data. Bioinformatics. 2016 Sep 15;32(18):2847-9.

49. Valero-Mora PM, Valero-Mora PM. Ggplot2: Elegant Graphics for Data Analysis. Journal of Statistical Software, Book of Reviews. 2010 Jul;35(1):1-3.