

# Interactions between culturable bacteria are predicted by individual species' growth

Einat Nestor<sup>1</sup>, Gal Toledano<sup>1</sup> and Jonathan Friedman<sup>2</sup>

<sup>1</sup>The Rachel and Selim Benin School of Computer Science and Engineering, Hebrew University, Jerusalem, Israel

<sup>2</sup>The institute of environmental science, Hebrew University, Rehovot, Israel

## Abstract

Predicting interspecies interactions is a key challenge in microbial ecology, as such interactions shape the composition and functioning of microbial communities. However, predicting microbial interactions is challenging since they can vary considerably depending on species' metabolic capabilities and environmental conditions. Here, we employ machine learning models to predict pairwise interactions between culturable bacteria based on their phylogeny, monoculture growth capabilities, and interactions with other species. We trained our models on one of the largest available pairwise interactions dataset containing over 7500 interactions between 20 species from 2 taxonomic groups that were cocultured in 40 different carbon environments. Our models accurately predicted both the sign (accuracy of 88%) and the strength of effects ( $R^2$  of 0.87) species had on each other's growth. Encouragingly, predictions with comparable accuracy could be made even when not relying on information about interactions with other species, which are often hard to measure. However, species' monoculture growth was essential to the model, as predictions based solely on species' phylogeny and inferred metabolic capabilities were significantly less accurate. These results bring us a step closer to a predictive understanding of microbial communities, which is essential for engineering beneficial microbial consortia.

## Introduction

Microbes are key participants in various processes, ranging from the health of humans<sup>1</sup>, animals and plants<sup>2</sup> to global biogeochemicals cycles<sup>3</sup>. The impact of microbes however, is usually not due to a single species but rather caused by diverse communities of interacting species<sup>4</sup>. Therefore, the mechanisms by which microbial species promote or hinder each other's growth has been studied extensively<sup>5</sup>. For example, negative effects can occur due to resource competition or secretion of antimicrobials<sup>6,7</sup>, whereas positive ones may occur due to cross-feeding<sup>8</sup> of metabolites, such as amino acids<sup>9</sup>.

Predicting interspecific interactions is necessary to understand a community's properties, as they are expected to be shaped by interactions within the community<sup>10-13</sup>. Indeed, pairwise interactions have been shown to be predictive of the structure and function of various simplified microbial communities<sup>14-18</sup>. However, it can be extremely challenging to directly measure all

pairwise interactions in a community or to infer them from sequencing data<sup>19–21</sup>. An alternative approach, which is likely essential for species-rich communities or ones comprised of fastidious species, is developing methodologies for predicting how microbes affect each other's growth.

Metabolic modeling and genome-based models have been commonly used to predict microbial interactions<sup>22–24</sup>. These approaches predict interactions by considering the overlap and complementarity between species' metabolic capabilities and/or their resource consumption and secretion<sup>25,26</sup>. These approaches are appealing since they rely solely on genomic information. However, their performance depends on the availability of well-annotated genomes, and they typically do not account for non-metabolic interaction modalities, such as the secretion of antibiotics or pH modifications<sup>27</sup>.

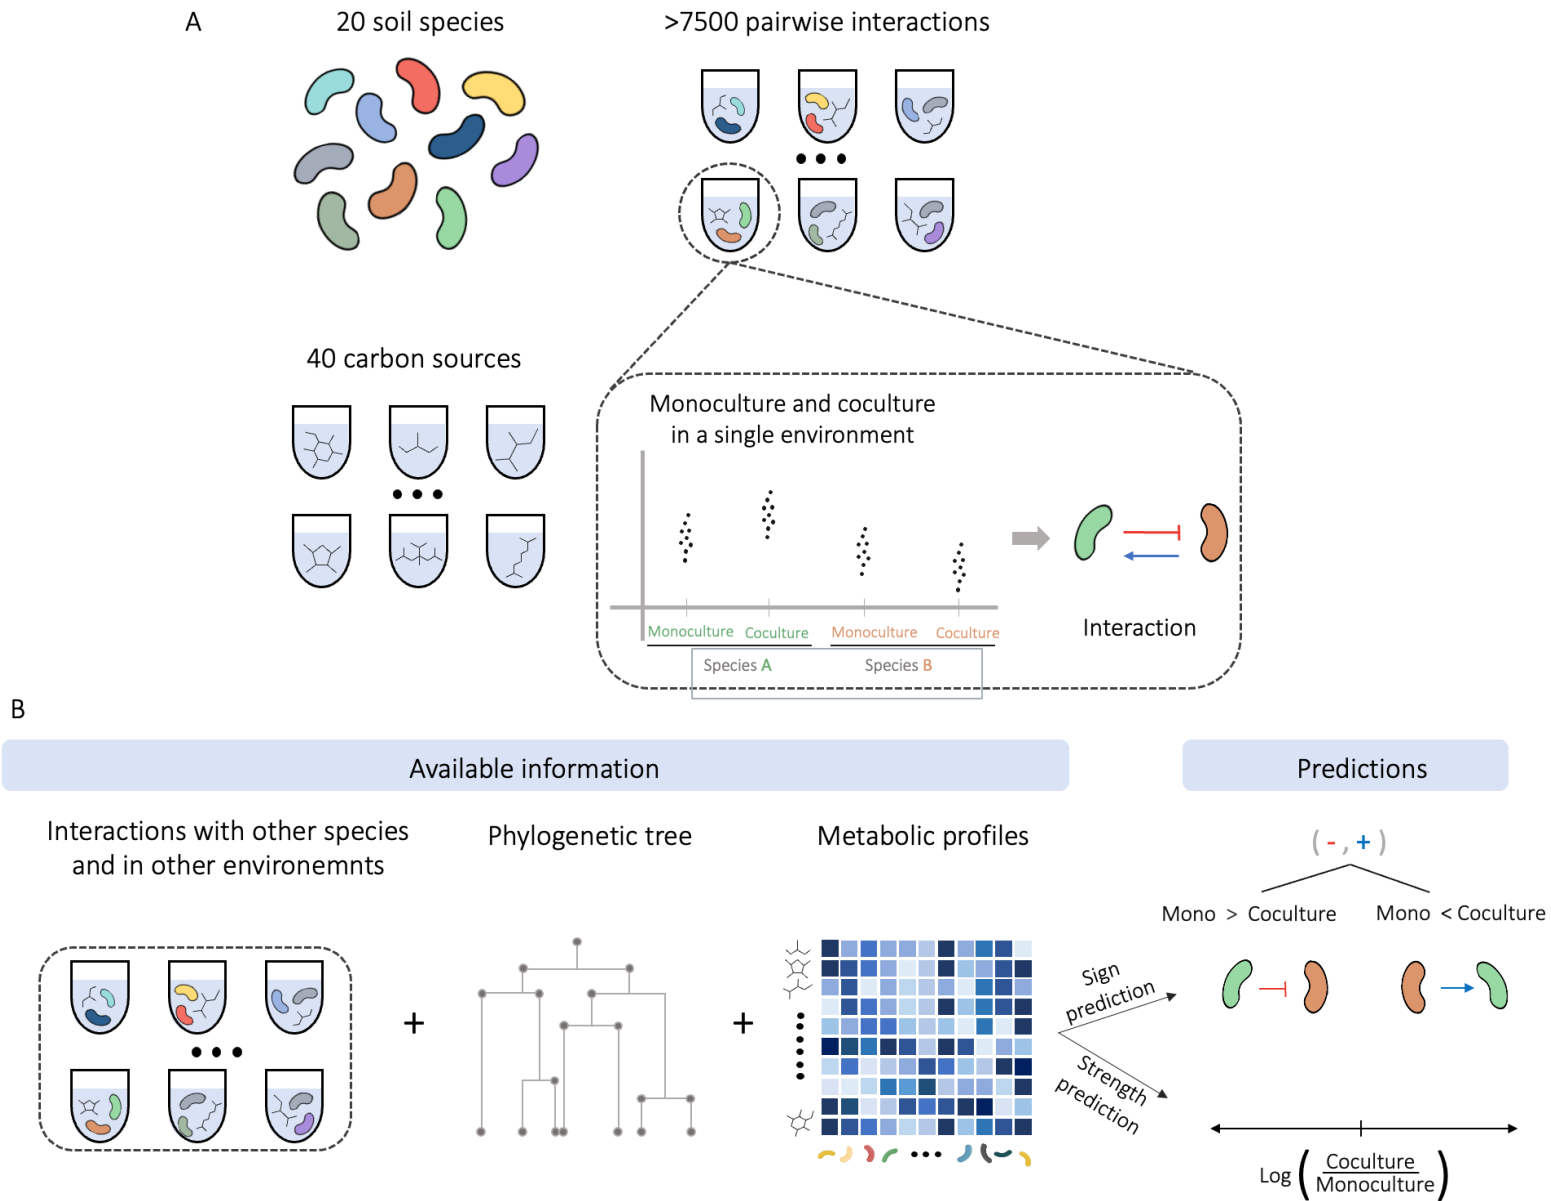
Another promising approach to interaction prediction is the use of machine learning models<sup>28</sup>. The use of supervised and unsupervised machine learning algorithms has increased in the past few years in many biological fields, including microbiology<sup>29,30</sup>. Previous works have managed to show that microbial community composition can be predicted using deep learning<sup>31,32</sup>. In addition, the use of supervised machine learning tools to accurately predict the sign of microbial interactions (positive or negative) based on genomic data and inferred metabolic pathways has recently been demonstrated<sup>33</sup>. While the latter results for bacterial interaction prediction are promising, they are restricted to engineered auxotrophic species, in silico simulations, and a handful of soil species in a single environment. Since interactions vary significantly between species and can drastically change across environments even between the same species, it is still not clear to what extent machine learning tools can predict interactions between non-engineered species across a range of nutrient conditions.

Here, we assess the ability of machine learning tools to predict microbial interactions using one of the largest datasets of experimentally-validated microbial interactions. This dataset contains all pairwise interactions among 20 different soil bacteria from 2 taxonomic groups that were cultured in 40 different media, each containing a single carbon source or a mixture of all carbon sources. Combined with phylogenetic information and phenotypic features, which were created from the dataset, machine learning models were able to accurately predict both the sign and the strength of pairwise interspecific interactions in this dataset.

## Results

In order to predict how species affect each other, we have used additional information, beyond the interspecific interactions, regarding the species' phylogeny and their monoculture yield in each of the 40 carbon environments. The growth of all species in monoculture and in coculture with each other species in each carbon environment was measured using the kChip combinatorial screening platform<sup>34</sup>. The one-way effect of one species on another in a given environment was quantified as the log ratio of the affected species' growth yield in coculture and in monoculture in that environment (see Methods). Information regarding each species' phylogeny and metabolic capabilities were included as features based on the species phylogenetic or monoculture growth profile similarity (represented as the first 2 principle

components, abbreviated as PCs, of the phylogenetic distance matrix for each species or as the first 4 PCs of the monoculture growth distance matrix, see Methods). We have first used this large dataset to train machine learning models to predict either the sign (positive/negative) or strength of one-way effects of one species on another's growth yield (Fig. 1).

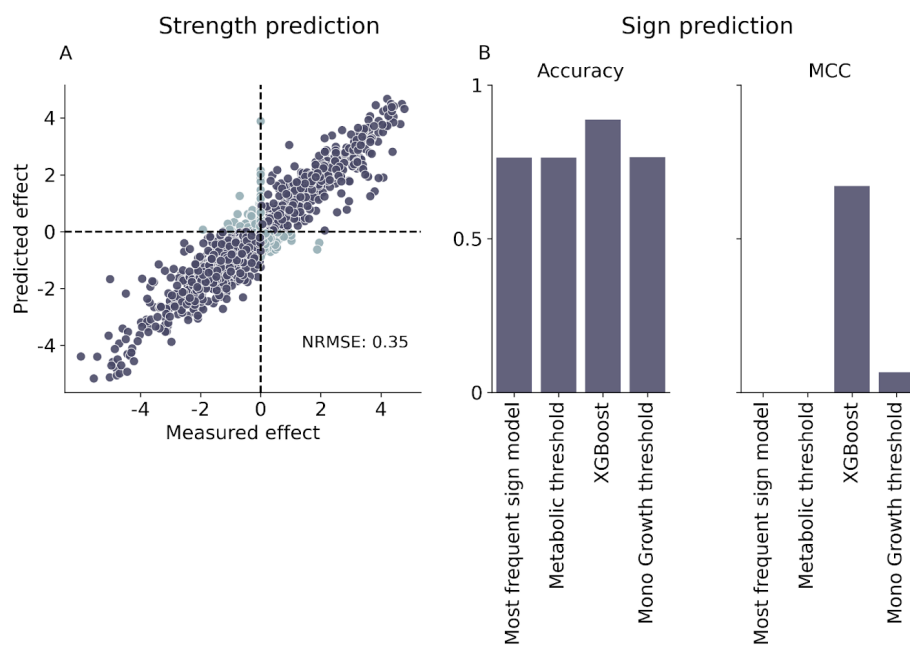


**Figure 1. An illustration of the data used to quantify and predict bacterial interactions. A.** All pairwise interactions among 20 different soil bacteria from 2 taxonomic groups, cocultured in 40 different carbon sources were measured. **B.** Information regarding species' phylogeny and metabolic capabilities were created from the dataset. Machine learning models were trained on the data, in order to predict the sign and strength of effects of one species on the growth of another.

### Machine learning algorithms predicted well both the sign and the strength of one-way growth yield effects

We evaluated the predictive ability of several machine learning algorithms and found that tree-based models performed best for predicting both effect sign and strength (XGBoost for both sign and strength, Fig. S1). The performance of these models was also superior to that of null models that always predict the most frequent sign/ average effect strength of train set, and threshold models that use a predefined threshold of a single feature (e.g. predict a negative effect if the metabolic distance between the interacting species is above a threshold, or the monoculture growth of the affected strain is above a threshold). These results confirm that machine learning models can predict both effect sign and strength better than models with a simple decision role.

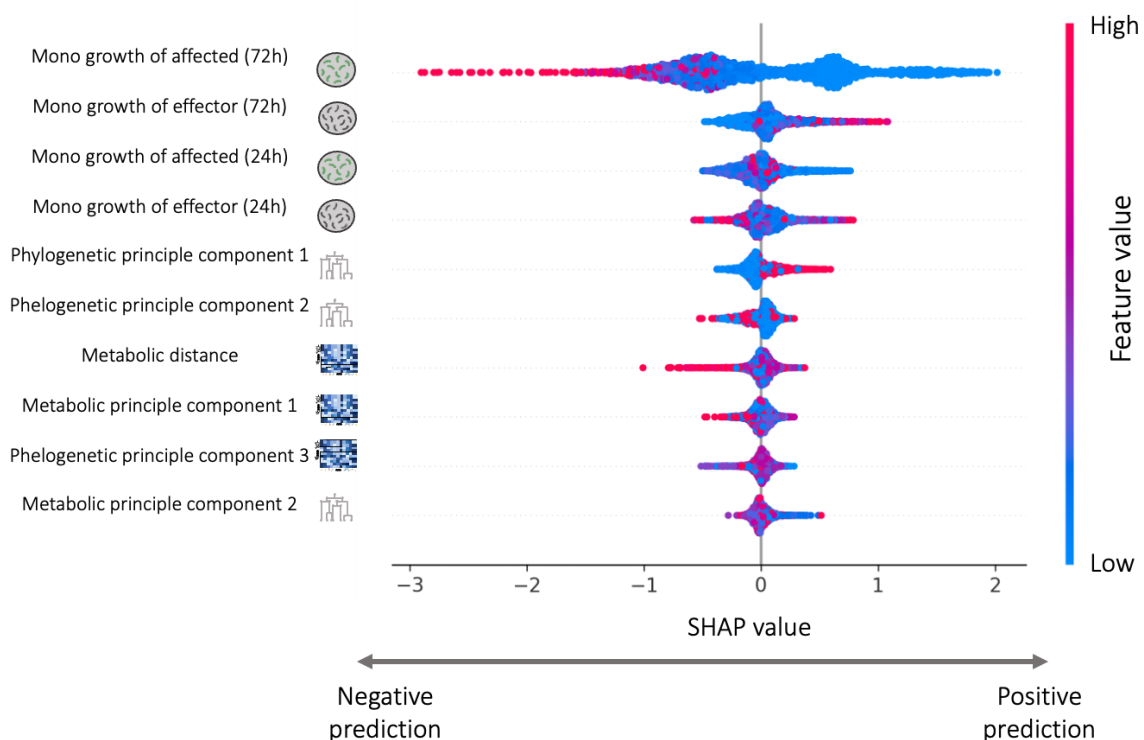
Tree-based models accurately predicted both the sign and strength of one-way effects. Quantitative predictions of effect strength achieved a normalized root-mean-square error (NRMSE) of 0.35 and  $R^2$  of 0.87 on the validation set (Fig. 2A). Qualitative predictions of effect sign had an out-of-sample accuracy of 0.88 as well as high Precision (0.7), Recall (0.81) and Mathews correlation coefficient (0.67), which accounts for the fact that our data is imbalanced with 73% negative effects (Fig. 2B, Fig. S2A). Most errors in effect sign prediction (3.5% false positives, 7.7% false negative) occurred for effects whose strength was close to 0 (Fig. S2B), indicating that the model was able to distinguish well between effects that were strongly negative or positive, but had difficulties in classifying weaker ones. In addition, while naive models achieved similar accuracy as the XGBoost, they performed very poorly in all other matrices (Fig. 2B, Fig. S2A), strengthening the conclusion that simple decision roles offered little predictive power.



**Figure 2. Tree-based machine learning models accurately predicted the sign and strength of one-way growth yield effects.** **A.** Measured effect vs predict effect for each species pair in the validation set for the best quantitative model, XGboost (Methods). Dark gray indicates samples with the same sign in actual effect and predicted effect while light blue indicates samples with different signs. **B.** Comparison of Accuracy and Matthews correlation coefficient for the best qualitative model (XGboost classifier), null model (assign common effect sign in train set) and 2 simple one feature decision role models: metabolic distance threshold model and monoculture growth of affected strain threshold model (Methods). Other metrics (recall and precision) are in S2.

### Monoculture growth yield is the most predictive feature

We further analyzed the contribution of each feature to the performance of the models using SHapley Additive exPlantations (SHAP), a game theory approach that measures the contribution of each feature to the total prediction of the model<sup>35,36</sup>. How well both species can grow in monoculture in the carbon environment in which they were interacting had the strongest influence on the prediction of both effect sign and strength (Fig. 3, Fig. S3). SHAP analysis indicates that higher monoculture growth yield of the affected species leads to a stronger negative contribution to the model's output. In other words, species that grow better in monoculture tend to be more negatively affected by the presence of additional species. This is consistent with previous findings that monoculture yields shape pairwise interactions<sup>37</sup>. Surprisingly, using information regarding species' predicted metabolic pathways, which were previously shown to be predictive of interactions, instead of information regarding monoculture growth did not improve the predictive ability over a model that only used the species' phylogeny (Fig. S4).



**Figure 3. Monoculture growth is the most important feature for predicting the strength of effect one species has on another's growth.** Top 10 most important features in the best performing strength prediction model (XGboost). High SHAP values indicate positive influence on the predicted effect, and low values indicate negative influence on the predicted effect. Y-axis represents the different features, sorted in descending order according to their contribution to the model. Each dot represents a simulation of the model with a single change at a single feature value. The different colors represent the value of the feature. Similar contribution of monoculture growth features appeared also in sign predictions (Fig. S3A).

While tree-based models offered improved predictive power compared to simpler models, they relied on having information regarding each species' monoculture growth and interactions with many other species in each carbon environment. Obtaining such information can be a laborious and challenging task, especially for species that are hard to culture under laboratory conditions. Therefore, we next studied how accurately we can predict the sign and strength of one-way effects when only partial information is available. To do so, we trained new models with only partial information regarding one of the species and compared the accuracy of prediction to those of the models trained using all the data.

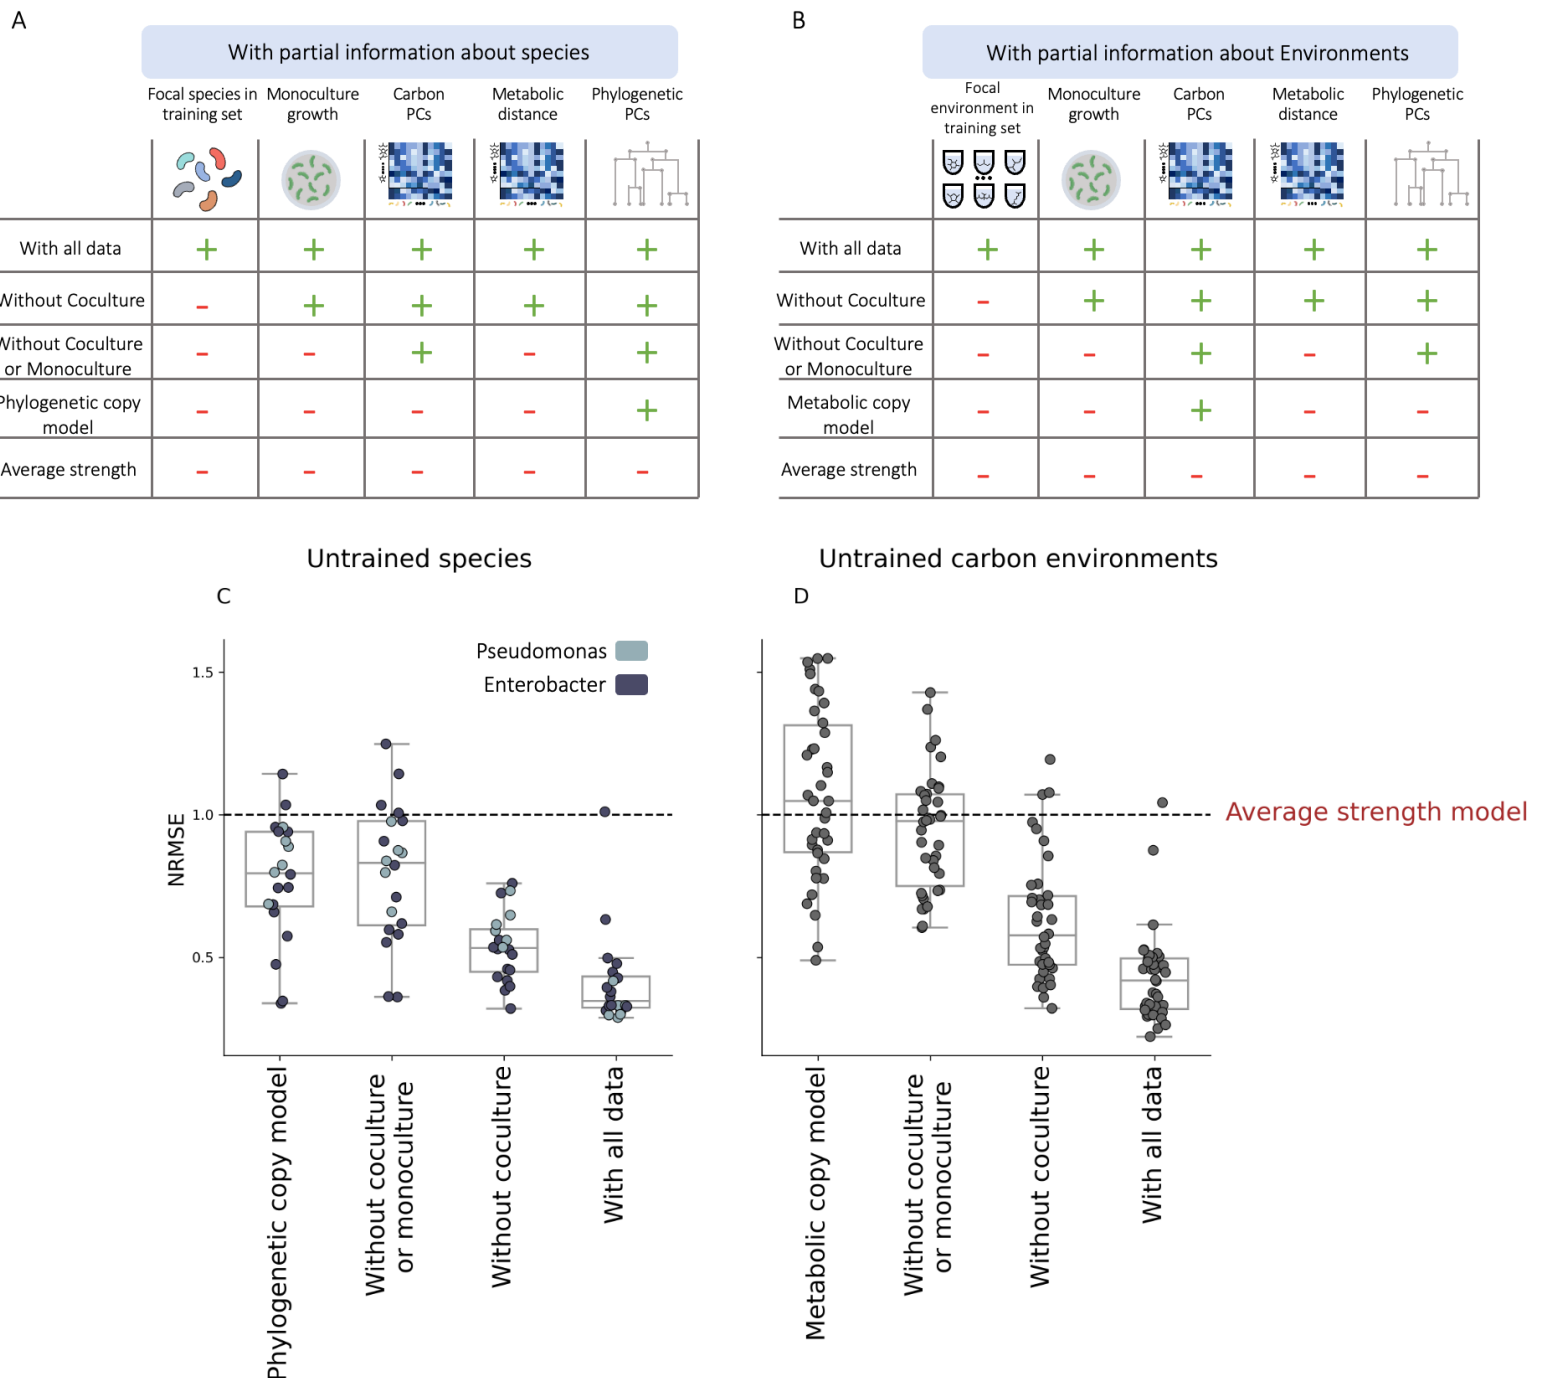
First, we evaluated our ability to predict interactions involving species for which we have monoculture growth data but no coculture data by removing a species from the training set. Next, we evaluated the accuracy of prediction when neither monoculture nor coculture data is available by removing a species from the training set and removing features related to monoculture growth from both the training and testing sets. In the latter case, predictions are based only on phylogenetic information. Lastly, we included a naive "phylogenetic copy" model where the sign or strength of the effect is assigned to be identical to those that involve the phylogenetically closest species in the same carbon environment (Fig. 4A, Methods).

The accuracy of predicting the sign and strength of one-way effects depended strongly on the availability of monoculture growth data, but not on coculture data (Figure 4.A,C). The lack of coculture data involving a given species increased the median prediction error (quantified using the NRMSE) by 0.18 (from 0.34 to 0.52), whereas removing monoculture data increased the median error by 0.3 (from 0.5 to 0.8). Moreover, when monoculture growth data is not available, prediction quality was similar to that of the simple "phylogenetic copy" model, which only requires the phylogenetic distance matrix (median RMSE values of 0.83 and 0.79).

In a similar way, we evaluate our ability to predict interactions that occur in a carbon environment for which we have only partial information. First, we removed a carbon environment from the training set. Next, we also removed features related to monoculture growth from both the training and testing sets. Lastly, we included a naive "metabolic copy" model where the sign or strength of effect is assigned to be identical to that of the same species in the metabolically closest carbon environment (Fig. 4B, Methods).



We again found that prediction accuracy depended more strongly on the availability of monoculture growth data than on coculture data (Fig. 4D), although overall predictions were less accurate. The lack of coculture data involving a given environment increased the median prediction error by 0.15 (from 0.42 to 0.57), whereas removing monoculture data increased the median error by 0.4 (from 0.57 to 0.9). The same pattern of improvement for sign predictions occurs among all tested metrics and for both species partial models and environment partial models (Fig. S5). These results indicate that if a species monoculture growth in a given carbon environment is known, growth effects involving that species can be well predicted given other species interactions in the same environment, or the same species' interactions in other environments.

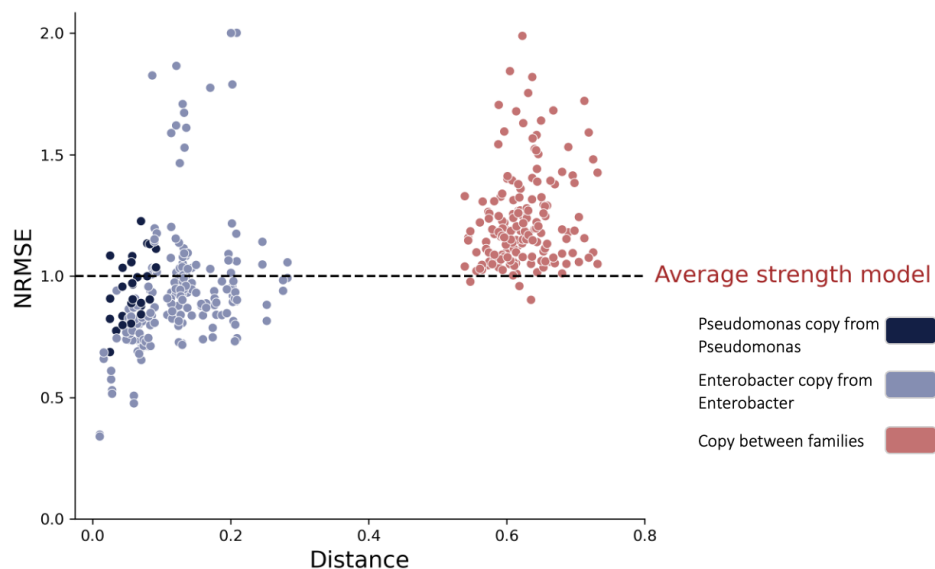


**Figure 4. The accuracy of predicting one-way effect depends strongly on the availability of monoculture growth data, but not on coculture data.**

**A.** A summary of the 4 compared models and the information used by each model. Partial information (left to right): Coculture, Monoculture (of both species), phylogenetic tree and type of model (machine learning or simple decision model). The models differ by the features they are trained on and their test set (see Methods). **B.** Comparison of NRMSE (normalized root mean square error) of several partial quantitative models, on “uncultured” strain. Each dot represents the NRMSE of a different strain that was excluded from the train set and was included in the test set. The models are the same as in A. Dashed line represents the performance of a null model (average effect in the train set). **C.** Comparison of NRMSE (normalized root mean square error) of several partial quantitative models, on “uncultured” carbon environment. Each dot represents the NRMSE of a different carbon source that was excluded from the train set and was included in the test set. The models are the same as in B. Dashed line represents the performance of a null model (average effect of train set). Qualitatively similar results were found also for predictions of effect sign (Fig. S5).

**Accuracy of “phylogenetic copy” model was higher for closely related species**

Since our best option for predicting interactions involving “uncultured” species (ones for which we have no monoculture or coculture data) was the simple “phylogenetic copy” model, we next examined how the phylogenetic distance from the “copied” species (for which interaction information is available) affects the prediction quality. As expected, the prediction accuracy and distance from “copied” species were significantly positively correlated (Fig.5; Pearson correlation coefficient 0.56, p-value <0.001). In other words, strains that are phylogenetically similar to the uncultured strain will be better predictors of the uncultured strain at various carbon sources whereas the greater the distance, the worse the predictions gets. However, poor prediction accuracy, lower than that achieved using the average effect strength, sometimes occurs even when copying interactions from species within the same family, and prediction accuracy varied between families. These results indicate that interactions tend to be conserved between closely related species, but the extent of conservation may vary between taxonomic groups.





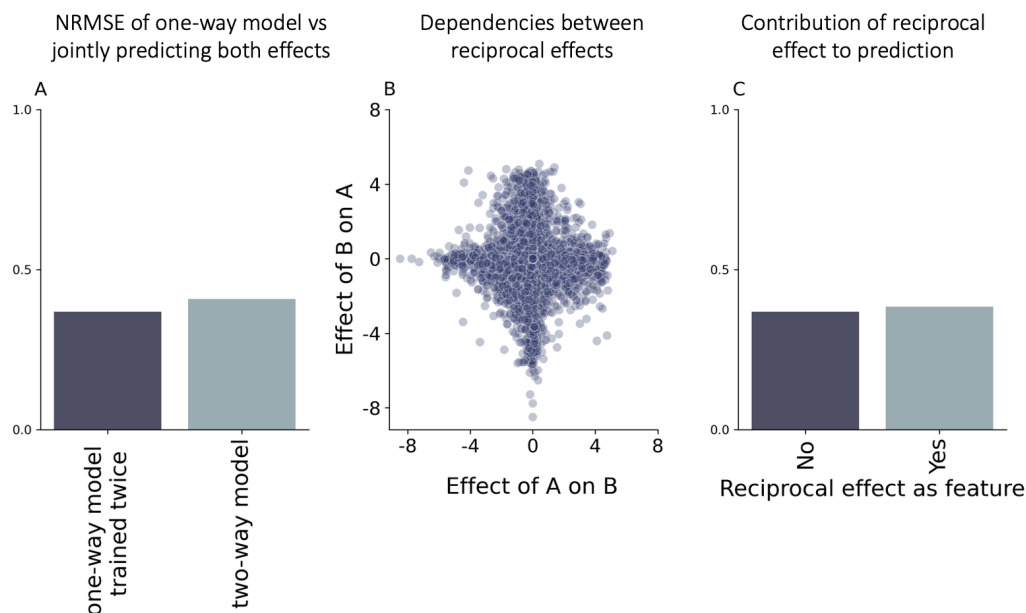
**Figure 5. A simple phylogenetic copy model can provide predictive power when based on closely related species.** Each dot represents a prediction for a single species, using one of the other 19 species as the copied species (interaction of the copied species in the same carbon source.) Different colors represent the distance category: the copied strain is from the same taxonomic family (dark blue, light blue) or from a different taxonomic family (red). Other metrics for sign predictions (Matthews correlation coefficient, Accuracy, recall and precision) are found in S6.

### **Combining one-way effect predictions is as accurate as jointly predicting two-way interactions**

Lastly, we studied how well we can predict two-way interactions that comprise both one-way effects of the interacting species on one another. We predicted two-way interactions using the same best-performing tree-based models that were used for predicting one-way effects (re-trained for multilabel output. See Methods). Similar to one-way effects, we quantified the accuracy of qualitative predictions of interaction type: competition (-, -), mutualism (+, +) and parasitism (-, +) and of quantitative predictions of interaction strength (Methods).

Surprisingly, jointly predicting two-way interactions was not more accurate than combining the independent predictions of two one-way effects (Fig. 6A,B Fig.S7). To better understand this finding, we quantified the dependence between reciprocal effects between a pair of species using Maximal Information Coefficient<sup>38</sup> (MIC) – a metric for capturing general dependencies between variables that ranges from 0 (independent) to 1 (fully dependent). Reciprocal effects between a pair of species were only weakly dependent on one another (MIC = 0.16), indicating that knowing how one species affects another isn't very predictive of the reciprocal effect.

In addition, we trained the same one-way strength model used to predict one-way effect, but with the reciprocal effect as an additional feature. Adding the reciprocal effect had little effect on prediction accuracy (NRMSE decrease of 0.02, Fig. 6C) and the reciprocal effect contributed little to predictions (Fig. S8). In other words, knowing the other species' effect doesn't add any helpful information, as this information is redundant when other features (monoculture growth yields, metabolic PC and phylogenetic PC) are available.



**Figure 6. Reciprocal effects are weakly dependent on one another and therefore do not improve the accuracy of interaction predictions. A.** NRMSE for using one-way strength model to predict two-way effects, and with multilabel strength model.

**B.** The relationship between the one-interaction effects. Each dot is a single interacting species, where the x-axis and the y-axis are the effects of one species on the other. **C.** NRMSE for one-way strength model without the reciprocal effect as a feature, and with the reciprocal effect as a feature.

## Discussion

Microbial interactions can help predict the properties of microbial communities, but are challenging to measure<sup>19–21</sup>. Here, we demonstrate that tree-based machine learning models can accurately predict the sign and, more importantly, the strength of bacterial interactions. These predictions were based on the species' phylogeny as well as on phenotypic features which are extracted from the monoculture growth yields of the species in various carbon courses.

The ability of the affected species to grow in monoculture in a given the carbon environment was the feature that contributes the most to prediction. Consistent with previous findings, species that grow well in monoculture are predicted to be more negatively affected by coculturing with other species, and to affect other species more positively. While prediction accuracy depended strongly on the availability of a species' monoculture growth data, it was less sensitive to the removal for coculture data. This is encouraging, since it indicates that the number of measurements allowing accurate interaction prediction scales linearly, rather than quadratically with the number of species.

In the absence of monoculture growth data, a simple phylogenetic copy model, which is intuitive and easy to create, offered some predictive power. In this model, interactions between a pair of species were predicted to be identical to those that occur between closely related species in the same environment. This indicates that bacterial interactions are to some extent phylogenetically conserved, at least within the two families analyzed here, and that known interactions may be informative regarding the interactions between other, closely related, species for which no growth data is available.

In contrast, predicting interactions in new carbon environments was significantly less accurate. Predicting that the interaction between a pair of species was identical to the interaction of the same pair in the most “similar” carbon environment was not accurate. This poor accuracy may be due to the fact that carbon sources are not clustered into distinct groups based on the species’ growth abilities (like species are clustered according to the phylogenetic tree; Fig. S8). More in-depth research is needed in order to best use information from one environment to make predictions regarding another environment, which may improve interaction predictions, especially in “new” carbon environments<sup>39</sup>.

Surprisingly, models that were trained using information regarding each species’ inferred metabolic pathways did not achieve higher prediction accuracy than models that used only phylogenetic information (Fig. S4). However, the metabolic pathways were inferred from the 16S sequences using picrust2<sup>40</sup>, rather than from well-annotated whole genome sequences. Therefore, it is possible that the addition of metabolic pathways that are constructed from whole-genome sequences will improve the performance of the models, improve prediction accuracy for uncultured species, and offer insights regarding the mechanistic basis underlying bacterial interactions.

Predicting pairwise interspecific interactions is crucial for understanding the structure, stability and function of microbial communities. Here, we demonstrate that tree-based machine learning models can be used for accurately predicting interactions of different species within the same taxonomic group or between different taxonomic groups, in a relatively large set of conditions (40 different carbon environments). Further work is needed in order to test the ability of this approach to predict interactions between more diverse taxonomic groups, and in more complex situations involving multiple species and nutrients. Being able to predict microbial interactions would put us one step closer to predicting the functionality of a microbial communities and to rationally microbiome engineering.

## Methods

### Data

The dataset contains over 7500 pairwise interactions involving 20 species from 2 taxonomic groups in 40 different carbon environments, as well as each species' monoculture growth yield in each carbon environment (see previous work<sup>37</sup>). Briefly, species' growth in mono- and coculture was assayed using the kChip platform - a high throughput nanodroplet-based platform for combinatorial screening<sup>41</sup>. These data were used to calculate the growth effect of one species on another as the log ratio of the affected species growth in coculture and in monoculture. Lastly, pairwise interactions given by both the effect of species A on B and the effect of species B on A<sup>37</sup>.

### Features creation

We created features representing species' phylogeny based on a previously published phylogenetic distance matrix of the 20 species<sup>37</sup>. We performed principal component analysis (PCA) on this matrix and used the first two principal components, which capture >95% of the variance, as features. Features that represent the carbon environments were based on the species' metabolic profiles, where the metabolic profile of each carbon environment is the monoculture growth yields of the 20 species. We performed PCA on the metabolic profile matrix and used the first four principal components, which capture >90% of the variance, as features. These features represent each carbon environment according to similarities in monoculture growth yields of the different species. In addition, we included a metabolic distance feature, which we calculated as the Euclidean distance between the monoculture-growth yields profiles of each pair of interacting species.

### Model training

First, the data was split into 2 groups - train and test set (80% and 20% respectively). Then, the hyperparameters of each model were tuned by performing 5-fold cross-validation on the train set and choosing the parameter values that resulted in the best performance (highest accuracy for qualitative predictions, lowest RMSE for quantitative predictions). For each hyperparameter, 2500 values were sampled uniformly from a given range, presented in supplementary Table 1. The models which were used for qualitative predictions are: Random forest classifier, Logistic regression, K nearest neighbors classifier, and XGBoost classifier. The models which were used for quantitative predictions are: Random forest regressor, XGboost, linear regression and K nearest neighbors regressor. All models were used from scikit-learn open-source package (python). Hypertuning was made using RandomGridSearch (scikit-learn 1.0.1).

### Naive models

In addition to machine learning models, we evaluated the performance of several simple prediction role models:

1. Null models: predict the effect sign to be the most frequent sign in the training set and the effect strength to be the average interaction strength in the training set.
2. Threshold models (for effect sign only): predict the effect sign based on whether the value of a single feature exceeds a threshold value. The threshold value was set to be

the one that maximized accuracy in the training set. Two threshold models were created - one based on the metabolic distance between a pair of species and a second model based on the monoculture growth yield of the affected species.

### **Models trained using partial information**

In order to evaluate our ability to predict interactions involving species or carbon environments for which only partial information is available, we created the following:

#### Partial information regarding a species

For each of the species a different test set was created containing only the species interactions, excluding all the interactions involving the species from the train set. For each species excluded from the training set, three machine learning models were trained using different sets of features:

1. Without coculture, but with monoculture growth measurements and phylogenetic features.
2. Without coculture or monoculture growth measurements, but with phylogenetic features.
3. Only with phylogenetic features.

Additionally, a simple decision role model was evaluated:

4. Phylogenetic copy model - copies the interaction (sign or strength) of the phylogenetically closest (according to the phylogenetic distance) strain in the same carbon environment, when interacting with the same partner.

Overall, 4\*20\*2 (4 types of models, 20 species and 2 types of prediction) models were trained and compared.

#### Partial information regarding a carbon environment

For each of the carbon environments a different test set was created containing only the interactions occurring in that environment, excluding all the interactions in that environment from the train set. For each carbon environment excluded from the training set, three machine learning models were trained using different sets of features:

1. Without coculture (in the specific environment), but with monoculture growth measurements and phylogenetic features.
2. Without coculture or monoculture growth measurements (in the specific environment), but with phylogenetic features.
3. Only with phylogenetic features.

Additionally, a simple decision role model was evaluated:

4. Metabolic distance model - copied the interaction (sign or strength) in the most similar carbon environment (according to the Euclidean distance of the environments' metabolic profiles).

As the monoculture growth yields were used for creating the metabolic representation of the different carbon environment, the metabolic representation of the carbon environment excluded from the training set was generated using the PCA of the other carbon environment. Overall,

4\*40\*2 (4 types of models, 40 carbon environments and 2 types of prediction) models were trained and compared.

### **Model performance evaluation**

The performance of models predicting effect sign was evaluated using Matthews correlation coefficient, which accounts for the fact that negative interactions are more frequent in our dataset (73%). The performance of models predicting effect strength was evaluated using normalized RMSE (NRMSE), defined as the RMSE divided by the standard deviation of the observed effects in the test set.

### **Two-way interactions prediction**

A two way interaction (between species A and B) is composed of a pair of reciprocal effects (Effect of B on A, Effect of A on B). There are two ways to predict two-way growth effect:

1. Train the one-way effect model and predict each of the two reciprocal effects independently.
2. Train a two-way model with multi label output (each prediction is in the form of [Effect of B on A, Effect of A on B]) and jointly predict the two-way interaction.



1. Cho, I. & Blaser, M. J. The human microbiome: at the interface of health and disease. *Nat. Rev. Genet.* **13**, 260–270 (2012).
2. Berendsen, R. L., Pieterse, C. M. J. & Bakker, P. A. H. M. The rhizosphere microbiome and plant health. *Trends Plant Sci.* **17**, 478–486 (2012).
3. Falkowski, P. G., Fenchel, T. & Delong, E. F. The microbial engines that drive Earth's biogeochemical cycles. *Science* **320**, 1034–1039 (2008).
4. Gill, S. R. *et al.* Metagenomic Analysis of the Human Distal Gut Microbiome. *Science* **312**, 1355–1359 (2006).
5. Hibbing, M. E., Fuqua, C., Parsek, M. R. & Peterson, S. B. Bacterial competition: surviving and thriving in the microbial jungle. *Nat. Rev. Microbiol.* **8**, 15–25 (2010).
6. Geesink, P. *et al.* Growth promotion and inhibition induced by interactions of groundwater bacteria. *FEMS Microbiol. Ecol.* **94**, fiy164 (2018).
7. Tyc, O. *et al.* Impact of interspecific interactions on antimicrobial activity among soil bacteria. *Front. Microbiol.* **5**, (2014).
8. Seth, E. C. & Taga, M. E. Nutrient cross-feeding in the microbial world. *Front. Microbiol.* **5**, (2014).
9. D'Souza, G. *et al.* Ecology and evolution of metabolic cross-feeding interactions in bacteria. *Nat. Prod. Rep.* **35**, 455–488 (2018).
10. Konopka, A., Lindemann, S. & Fredrickson, J. Dynamics in microbial communities: unraveling mechanisms to identify principles. *ISME J.* **9**, 1488–1495 (2015).
11. Barbier, M., Arnoldi, J.-F., Bunin, G. & Loreau, M. Generic assembly patterns in complex ecological communities. *Proc. Natl. Acad. Sci.* **115**, 2156–2161 (2018).
12. Qian, J. J. & Akçay, E. The balance of interaction types determines the assembly and stability of ecological communities. *Nat. Ecol. Evol.* **4**, 356–365 (2020).

13. Coyte, K. Z., Schluter, J. & Foster, K. R. The ecology of the microbiome: Networks, competition, and stability. *Science* **350**, 663–666 (2015).
14. Friedman, J., Higgins, L. M. & Gore, J. Community structure follows simple assembly rules in microbial microcosms. *Nat. Ecol. Evol.* **1**, 1–7 (2017).
15. Ortiz, A., Vega, N. M., Ratzke, C. & Gore, J. Interspecies bacterial competition regulates community assembly in the *C. elegans* intestine. *ISME J.* **15**, 2131–2145 (2021).
16. Gould, A. L. *et al.* Microbiome interactions shape host fitness. *Proc. Natl. Acad. Sci.* **115**, E11951–E11960 (2018).
17. Venturelli, O. S. *et al.* Deciphering microbial interactions in synthetic human gut microbiome communities. *Mol. Syst. Biol.* **14**, e8157 (2018).
18. Bucci, V. *et al.* MDSINE: Microbial Dynamical Systems INference Engine for microbiome time-series analyses. *Genome Biol.* **17**, 121 (2016).
19. Carr, A., Diener, C., Baliga, N. S. & Gibbons, S. M. Use and abuse of correlation analyses in microbial ecology. *ISME J.* **13**, 2647–2655 (2019).
20. Cao, H.-T., Gibson, T. E., Bashan, A. & Liu, Y.-Y. Inferring human microbial dynamics from temporal metagenomics data: Pitfalls and lessons. *BioEssays News Rev. Mol. Cell. Dev. Biol.* **39**, (2017).
21. Hirano, H. & Takemoto, K. Difficulty in inferring microbial community structure based on co-occurrence network approaches. *BMC Bioinformatics* **20**, 329 (2019).
22. Shoaie, S. *et al.* Understanding the interactions between bacteria in the human gut through metabolic modeling. *Sci. Rep.* **3**, 2532 (2013).
23. Lara, E. G., van der Windt, I., Molenaar, D., de Vos, M. G. J. & Melkonian, C. Using Functional Annotations to Study Pairwise Interactions in Urinary Tract Infection Communities. *Genes* **12**, 1221 (2021).
24. Zoccarato, L., Sher, D., Miki, T., Segrè, D. & Grossart, H.-P. A comparative whole-genome approach identifies bacterial traits for marine microbial interactions. *Commun.*

- Biol.* **5**, 1–13 (2022).
25. Gu, C., Kim, G. B., Kim, W. J., Kim, H. U. & Lee, S. Y. Current status and applications of genome-scale metabolic models. *Genome Biol.* **20**, 121 (2019).
  26. Kim, T. Y., Sohn, S. B., Kim, Y. B., Kim, W. J. & Lee, S. Y. Recent advances in reconstruction and applications of genome-scale metabolic models. *Curr. Opin. Biotechnol.* **23**, 617–623 (2012).
  27. Ankrah, N. Y. D. *et al.* Enhancing Microbiome Research through Genome-Scale Metabolic Modeling. *mSystems* **6**, e00599-21 (2021).
  28. DiMucci, D., Kon, M. & Segrè, D. Machine Learning Reveals Missing Edges and Putative Interaction Mechanisms in Microbial Ecosystem Networks. *mSystems* **3**, e00181-18 (2018).
  29. Qu, K., Guo, F., Liu, X., Lin, Y. & Zou, Q. Application of Machine Learning in Microbiology. *Front. Microbiol.* **10**, (2019).
  30. Lopatkin, A. J. & Collins, J. J. Predictive biology: modelling, understanding and harnessing microbial complexity. *Nat. Rev. Microbiol.* **18**, 507–520 (2020).
  31. Baranwal, M. *et al.* Deep Learning Enables Design of Multifunctional Synthetic Human Gut Microbiome Dynamics. bioRxiv (2021).
  32. Michel-Mata, S., Wang, X.-W., Liu, Y.-Y. & Angulo, M. T. Predicting microbiome compositions from species assemblages through deep learning. *iMeta* **1**, e3 (2022).
  33. DiMucci, D., Kon, M. & Segrè, D. Machine Learning Reveals Missing Edges and Putative Interaction Mechanisms in Microbial Ecosystem Networks. *mSystems* **3**, e00181-18.
  34. Kehe, J. *et al.* Positive interactions are common among culturable bacteria. *Sci. Adv.* **7**, eabi7159 (2021).
  35. Lundberg, S. M. & Lee, S.-I. A Unified Approach to Interpreting Model Predictions. in *Advances in Neural Information Processing Systems* vol. 30 (Curran Associates, Inc., 2017).
  36. Lundberg, S. M. *et al.* From local explanations to global understanding with explainable

- AI for trees. *Nat. Mach. Intell.* **2**, 56–67 (2020).
37. Kehe, J. *et al.* Positive interactions are common among culturable bacteria. *Sci. Adv.* **7**, eabi7159.
38. Reshef, D. N. *et al.* Detecting Novel Associations in Large Data Sets. *Science* **334**, 1518–1524 (2011).
39. Estrela, S., Sanchez-Gorostiaga, A., Vila, J. C. & Sanchez, A. Nutrient dominance governs the assembly of microbial communities in mixed nutrient environments. *eLife* **10**, e65948 (2021).
40. Douglas, G. M. *et al.* PICRUSt2 for prediction of metagenome functions. *Nat. Biotechnol.* **38**, 685–688 (2020).
41. Kehe, J. *et al.* Massively parallel screening of synthetic microbial communities. *Proc. Natl. Acad. Sci. U. S. A.* **116**, 12804–12809 (2019).

