Predicting metabolic response to dietary intervention using deep learning

Tong Wang[1], Hannah D. Holscher[2,3], Sergei Maslov[4,3], Frank B. Hu[1,5,6], Scott T. Weiss[1], Yang-Yu Liu[1,3,*]

[1]*Channing Division of Network Medicine, Department of Medicine, Brigham and Women's Hospital, Harvard Medical School, Boston, MA 02115, USA.*
[2]*Department of Food Science and Human Nutrition, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA.*
[3]*Center for Artificial Intelligence and Modeling, The Carl R. Woese Institute for Genomic Biology, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA.*
[4]*Department of Bioengineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA.*
[5]*Department of Nutrition, Harvard T.H. Chan School of Public Health, Boston, MA 02115, USA*
[6]*Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, MA 02115, USA*
* Correspondence: yyl@channing.harvard.edu

Abstract

Due to highly personalized biological and lifestyle characteristics, different individuals may have different metabolic responses to specific foods and nutrients. In particular, the gut microbiota, a collection of trillions of microorganisms living in our gastrointestinal tract, is highly personalized and plays a key role in our metabolic responses to foods and nutrients. Accurately predicting metabolic responses to dietary interventions based on individuals' gut microbial compositions holds great promise for precision nutrition. Existing prediction methods are typically limited to traditional machine learning models. Deep learning methods dedicated to such tasks are still lacking. Here we develop a new method McMLP (**M**etabolic response predictor using **c**oupled **M**ulti**l**ayer **P**erceptrons) to fill in this gap. We provide clear evidence that McMLP outperforms existing methods on both synthetic data generated by the microbial consumer-resource model and real data obtained from six dietary intervention studies. Furthermore, we perform sensitivity analysis of McMLP to infer the tripartite food-microbe-metabolite interactions, which are then validated using the ground-truth (or literature evidence) for synthetic (or real) data, respectively. The presented tool has the potential to inform the design of microbiota-based personalized dietary strategies to achieve precision nutrition.

Introduction

Precision nutrition aims to provide personalized dietary recommendations based on an individual's unique biological and lifestyle characteristics such as genetics, gut microbiota, metabolomic profiles, and anthropometric data[1,2]. In addition to the design and implementation of large-scale clinical studies, one of the critical components for achieving precision nutrition is the development of predictive models that incorporate diverse individual data types to achieve an accurate prediction of metabolomic profiles following dietary changes[1–3]. However, existing models are limited to traditional machine learning methods such as Random Forest (RF)[4,5] and Gradient-Boosting Regressor (GBR)[3]. Deep learning techniques have not been leveraged to predict metabolic responses for precision nutrition.

Among the biological characteristics relevant for precision nutrition, the gut microbiota is an important factor that explains a large fraction of individual metabolic responses among populations[4–7]. Indeed, the human gut microbiota produces many metabolites through the microbial metabolism of nondigested food components such as dietary fibers, which are prevalent in grains, vegetables and fruits[8]. Increasingly, microbial metabolites have been shown to impact host health[9–13]. For example, short-chain fatty acids (SCFAs) are metabolites produced by intestinal microbes through anaerobic fermentation of indigestible polysaccharides such as dietary fiber and resistant starch[9,10]. SCFA concentrations have been linked to the regulation of immune cell function, gut-brain communication[16], and cardiovascular diseases[17,18]. Among the SCFAs, butyrate has been shown to be negatively correlated with pro-inflammatory cytokines[19,20]. Hence, a high level of butyrate from the gut microbiota is believed to be beneficial due to its anti-inflammatory effects[19–21]. Boosting the levels of health-beneficial metabolites by modulating the gut microbiota appears to be a promising approach to improve host health[22–24].

One possible way to modulate the gut microbiota is through dietary interventions[6]. Gut microbial composition is affected by the diet[6,25–28]. As a result, microbiota-targeted dietary interventions have been proposed to modulate the gut microbiota to increase the production of metabolites beneficial to the host[29–31]. Recently, there has been a growing trend to exploit the tripartite relationship between food/nutrition, gut microbiota, and microbiota-derived metabolites to provide better dietary advice for each individual[3–5,28–32]. Indeed, accurate prediction of personalized metabolic responses to foods and nutrients based on our gut microbiota holds great promise for precision nutrition[33].

Many dietary intervention studies have attempted to investigate the relationship
between diet and microbial metabolism of the gut microbiota[27,28,32,34]. However, most of these
studies only analyzed correlations between dietary treatments, microbes or metabolites. A few
studies have used different analytic approaches to predict postprandial metabolic responses
of markers such as blood glucose[3,4] and immune markers[4,34]. However, the personalized
prediction of how important markers such as SCFAs and bile acids respond to long-term
dietary interventions is under-investigated (Fig. 1).
Herein, we leveraged data from randomized, controlled dietary intervention
studies[27,28,35–38] and developed a deep-learning method: **M**etabolic response predictor using
coupled **M**ulti**l**ayer **P**erceptrons (McMLP) to predict post-dietary intervention metabolite
concentrations based on pre-dietary intervention microbial composition. We first proposed a
microbial consumer-resource model with cross-feeding interactions to simulate the dietary
intervention process and generate synthetic data to validate McMLP. We found that McMLP
outperforms existing methods (RF and GBR), especially when the training sample size is small.
We then applied all methods to real data from six dietary intervention studies[27,28,35–38], finding
that the predictive power of McMLP is higher than existing methods. Finally, based on the well-
trained McMLP, we performed sensitivity analysis to infer the tripartite food-microbe-metabolite
relationship. This helps us identify key microbes that serve as both strong consumers of the
intervened food and strong producers of the beneficial metabolite we would like to boost. We
presented some literature evidence that supports the tripartite food-microbe-metabolite
relationship inferred by McMLP.
89

Results

Overview of McMLP

Our aim is to predict the post-dietary intervention (or "endpoint") metabolite concentrations in
fecal or blood samples based on the pre-dietary intervention (or "baseline") microbial
composition, metabolome data, and the dietary intervention strategy. This is conceptually
different from existing studies on the inference of metabolomic profiles from microbial
compositions measured at the same time[39–42]. We hypothesized that in order to accurately
predict post-dietary intervention metabolomic profiles, we first need to capture how microbial
composition changes from the baseline to the endpoint. This is because metabolomic profiles
reflect the microbial metabolism of a community[7,43].

To test our hypothesis, we proposed McMLP, which consists of two steps : (step-1) use
the baseline microbiota and metabolome data and the dietary intervention strategy to predict
the endpoint microbial composition; and (step-2) use the predicted endpoint microbial
composition, the baseline metabolome data, and the dietary intervention strategy to predict the
endpoint metabolomic profile (Fig. 2a; Supplementary Fig. 1a). For each step, we used a
multilayer perceptron (MLP) with Rectified Linear Unit (ReLu) as the activation function to
perform the prediction. We emphasize that, in principle, one can just use one MLP to directly
predict endpoint metabolomic profiles based on baseline microbiota/metabolome data and the
dietary intervention strategy (Supplementary Fig. 1b). Later, we confirmed that this one-step
strategy has worse predictive power than our two-step strategy.

From a practical standpoint, our goal is to predict an individual's metabolic response to
a potential dietary intervention to facilitate precision nutrition. To achieve this goal, we feed the
baseline microbiota and metabolome profiles of this individual and the potential dietary
intervention strategy to a well-trained McMLP to predict the endpoint metabolome profile. Note
that in this application (or test) stage, because the dietary intervention is a thought experiment,
no real endpoint data is available. The first MLP in McMLP will predict the endpoint microbiota
profile, which will be fed into the second MLP to predict the endpoint metabolome profile.

During the training stage of McMLP, we need to collect not only baseline microbiota
and metabolome profiles of different individuals, but also perform dietary interventions to collect
actual endpoint microbiota and metabolome profiles. We emphasize that the actual endpoint
microbiota data will only be used to train the first MLP (Fig. 2b). It shall not be used to train the
second MLP. This is because we need to keep the consistency between the training and
application (or test) stages. After all, during the application stage, it is the predicted endpoint
microbiome profile that will be fed into the second MLP, and the actual endpoint microbiome
profile does not exist at all.

Instead of fine-tuning hyperparameters such as the number of layers N_l and the hidden
layer dimension N_h for MLP, we overparameterized MLP by using a large and fixed number of
layers N_l and hidden layer dimension N_h ($N_\mathrm{l} = 6$ and $N_\mathrm{h} = 2048$). The overparameterized
machine learning methods, especially deep learning models, yield better performance due to
their high capacity (i.e., more model parameters). In fact, the high-capacity models can be even
simpler due to smoother function approximation and thus less likely to overfit[44].

To illustrate the prediction task, we used a hypothetical example comprising $N_\mathrm{s}(= 5)$
microbial species, $N_\mathrm{d}(= 3)$ dietary resources being intervened, $N_\mathrm{m}(= 6)$ metabolites, and 7
samples (Fig. 2b,c). We will use both the baseline data and the dietary intervention strategy as

inputs for McMLP (Fig. 2a). 5 samples are used as the training set (Fig. 2b) and the remaining

2 samples form the test set (Fig. 2c). To evaluate the regression performance, we employed

three metrics based on the Spearman correlation coefficient (SCC) ρ between the predicted

and true values of the concentration of one metabolite across all samples: (1) $\bar{\rho}$: the mean

SCC, (2) $f_{\rho > 0.5}$: the fraction of metabolites with ρ greater than 0.5, and (3) $\bar{\rho}_5$: the mean SCC

of the top-5 best-predicted metabolites.

McMLP generates superior performance over existing methods on synthetic data

To validate the predictive power of McMLP, we applied it to synthetic data generated from the

Microbial Consumer-Resource Model (MiCRM) which considers microbial interactions through

both nutrient competition and metabolic cross-feeding[45]. We adapted MiCRM to simulate the

dietary intervention. For simplicity, we considered 20 food resources, 20 microbes, and 20

metabolites in the modeling. Also, we assumed that food resources can only be consumed

while metabolites can be either consumed or produced. Prior to the dietary intervention, one

food resource (referred to as "food resource #1") was not introduced, while the remaining 19

food resources were supplied. Dietary intervention was simulated by adding food resource #1

at a specific "dose" to microbial communities composed of surviving species before the dietary

intervention and calculating the new ecological steady state. Here, the "dose" is defined as the

ratio between the concentration of the intervened food resource and that of other food

resources.   We split the synthetic data (with 250 samples) with 80/20 ratio five times to

generate five train-test pairs that can be used to reflect the variation in predictive performance.

Details on model simulation and synthetic data generation can be found in the Supplementary

Information.

We compared the performance of McMLP with two classical methods (GBR: Gradient-

Boosting Regressor[3]; RF: Random Forest[4,5]) in the prediction task defined in Fig. 2. For each

method, we considered two sets of input variables: (1) without baseline metabolomic profiles

(denoted as "w/o b" hereafter) and (2) with baseline metabolomic profiles (denoted as "w/ b"

hereafter).

We first used the three metrics ($\bar{\rho}$, $f_{\rho > 0.5}$, $\bar{\rho}_5$) to benchmark the predictive performance

of the different methods on synthetic data with 50 training samples and an intervention dose of

3. We found that McMLP generated the best performance (Figs. 3a1-a3), especially when

baseline metabolomic profiles were included in the input. When we predict without baseline

metabolomic profiles, McMLP is clearly better than RF and GBR (McMLP yields the highest $\bar{\rho}$

of 0.399 ± 0.014, the highest $f_{\rho>0.5}$ of 0.200 ± 0.062, and the highest $\bar{\rho}_5$ of 0.538 ± 0.014; the

standard error is used to measure the variation in performance metrics across 5 train-test

splits). Including baseline metabolomic profiles in the input significantly improves the

performance of all methods, with McMLP still being the best (which yields the highest $\bar{\rho}$ of

0.613 ± 0.012, the highest $f_{\rho>0.5}$ of 0.860 ± 0.036, and highest $\bar{\rho}_5$ of 0.730 ± 0.010). We also

tried to introduce 5 food resources during the dietary intervention (instead of 1 previously; see

Supplemental Information for details) and found that the performance of McMLP is still superior

to other methods when the dietary intervention strategy is more complex (Supplementary Fig.

2).

We further examined the effect of training sample size on model performance. While

maintaining the same 50-sample test set used previously, we found that all performance

metrics for all methods improved as the training sample size increased (Fig. 3b1-b3). More

importantly, we found that the performance of McMLP is better than RF and GBR at small

training sample sizes (20 or 50) and is close to RF and GBR at large training sample sizes

(>50). This demonstrates the superior performance of McMLP with a limited number of samples,

contrary to the traditional notion that deep learning methods tend to overfit at small sample

sizes[46].

We finally examined the effect of intervention dose on model performance. By varying

the concentration of the intervened food resource in MiCRM, we generated synthetic data with

different intervention doses and subsequently trained all ML methods on them with 200 training

samples. We found that the performance gap between methods using and not using baseline

metabolomic profiles narrows as the intervention dose increases (Fig. 3c1-c3). We believe this

is because a larger intervention dose significantly changes the endpoint metabolomic profile

away from its baseline level, rendering the baseline metabolomic profile less useful.

McMLP accurately predicts metabolic responses on real human gut microbiota

data

After validating McMLP using synthetic data, we analyzed real data from six dietary intervention

studies to see if its performance on real data was consistently better than existing methods.

The first dataset we collected was from a study investigating how avocado consumption alters

gut microbial compositions and concentrations of metabolites such as SCFAs and bile acids[28].

In this study all participants were divided into two groups based on the food components of the

meals provided: (1) avocado group: 175 g (men) or 140 g (women) of avocado was provided

as part of a meal once a day for 12 weeks and (2) control group: no avocado was included in

their control meal[28]. Baseline (i.e., before the dietary intervention) and endpoint (i.e., during

week 12 of the intervention) microbial compositions and concentrations of SCFAs and bile

acids were quantified. The dataset is unique due to its relatively large sample size (66 for both

avocado and control groups)[28] compared to other dietary intervention studies[27,32,34].

Because the amount of avocado consumed by participants in the avocado group was

very similar and participants in the control group barely consume avocado, for simplicity, we

encoded the participant's dietary intervention in McMLP and other methods as a binary variable

in the input (green icons/symbols representing diets in Fig. 2) whose value equals 1 or 0 if the

participant is in the avocado or control group, respectively. Note that in this study the

concentrations of SCFAs and bile acids were obtained from two separate targeted

metabolomic assays. Hence, we separated the concentration prediction of SCFAs and bile

acids to compare the predictability of the two metabolite classes. We found that for the

concentration prediction of both SCFAs and bile acids, McMLP with the baseline metabolomic

profiles consistently produces the best performance (Fig. 4a1-a3, b1-b3). Interestingly, the

inclusion of baseline metabolomic profiles in the input of McMLP helps more with the prediction

of bile acid concentrations than with the prediction of SCFA concentrations ($\bar{\rho}$ increases from

0.226 to 0.396 for bile acids when metabolomic profiles are included; $\bar{\rho}$ increases from 0.302

to 0.385 for SCFAs when metabolomic profiles are included). A potential explanation is that

the correlation of SCFA concentrations between baseline and endpoint samples is weaker than

that of bile acids (Supplementary Fig. 3).

We checked the predictive performance of the one-step strategy (Supplementary Fig.

1b), finding that it is not as good as that of McMLP (Supplementary Fig. 4). We also compared

McMLP with the state-of-art method of predicting metabolomic profiles from microbial

compositions measured at the same time --- mNODE[42], finding that it has worse performance

than McMLP (Supplementary Fig. 5). The worse performance of mNODE is likely due to the

fact that it is not dedicated to predicting metabolomic profiles at different time points. More

technical reasons can be found in the Supplementary Information.

We extended the method comparison to five additional datasets from independent

dietary studies investigating how microbiota compositions and metabolomic profiles were

influenced by adding grains[35], walnuts[27], almonds[36], broccoli[37], and high-fiber or fermented

foods[38] (see Table 1 and Methods section for details of the studies). Each participant's dietary

intake was similarly encoded as either a binary variable or a vector whose value is proportional

to the consumed amount of the added dietary component, depending on the complexity of the

dietary intervention. Further details of the data processing and model architecture setup can

be found in the Supplementary Information. As shown in Fig. 4, McMLP consistently produces

the best performance across all datasets. The relatively poor performance of all methods on

the data from the study that investigated fibers and fermented foods[38] is likely due to the fact

that a variety of foods within the fiber and fermented foods categories were consumed by the

participants at will, while other studies were complete feeding trials[38].

Inferring the tripartite food-microbe-metabolite relationship

It has been previously shown that an individual's metabolic response depends on her/his gut

microbial composition[7,43]. If we want to introduce a new dietary resource to boost the

concentration of a health-beneficial metabolite mediated by gut microbes, we need "key"

microbial species that meet two criteria: (1) the species can consume one or more nutrient

components in the introduced food resource; (2) the species can produce the metabolite we

want to boost. If either criterion is not met, it is difficult to boost the metabolite concentration

via this dietary intervention. We aim to identify these "key" species that satisfy both criteria by

revealing the food-microbe consumption and microbe-metabolite production patterns, which

can be summarized in a tripartite food-microbe-metabolite graph (Supplementary Fig. 6). To

achieve this, we performed sensitivity analysis of McMLP. In particular, we interpreted a

potential relationship between an input variable x and an output variable y by perturbing x by

a small amount (denoted as Δx) and then measuring the response of y (denoted as Δy).

Following the notion of sensitivity in engineering sciences, we defined sensitivity $s = \frac{\Delta y}{\Delta x}$ and

used its sign (positive/negative) to reflect whether y changes in the same/opposite direction as

x. Details of this calculation can be found in the Methods section and in our previous study[42].

We calculated sensitivities for step-1 (and step-2) in McMLP to infer potential food-

microbe consumption (and microbe-metabolite production) interactions, respectively (Fig. 5a).

Specifically, in step-1, we perturbed the amount of food resource α and measured the change

in the relative abundance of species i. The sensitivity of species i to food resource α is $s_{i\alpha} =$

$\frac{\Delta y_i}{\Delta x_\alpha}$ and its sign can be used to reflect the interaction between species i and food resource α.

$s_{i\alpha} > 0$, indicates that species i can consume some nutrient components of food resource α.

Similarly, for step-2, we define the sensitivity of metabolite β to species i as $s_{\beta i} = \frac{\Delta y_\beta}{\Delta x_i}$. The

positive sensitivity, $s_{\beta i} > 0$, reveals potential production of the metabolite β by species i.

We first evaluated our sensitivity method on the synthetic data for which we know the
ground truth of food-microbe consumption and microbe-metabolite production interactions. We
found that the inferred sensitivity values for all food-microbe and microbe-metabolite pairs (Fig.
5b) have a zero-nonzero pattern very similar to the ground-truth consumption and production
rates assigned in MiCRM (Fig. 5c). We chose zero as the sensitivity threshold and kept only
positive values for food-microbe pairs (green cells in Fig. 5b&c) and for microbe-metabolite
pairs (red cells in Fig. 5b&c) to explore consumption and production interactions respectively.
To statistically verify the agreement between ground-truth interactions and inferred interactions
based on sensitivity values, we computed the AUROC (Area Under the Receiver Operating
Characteristic curve) based on the overlap between true and predicted interactions when the
classification threshold is varied. More specifically, for each classification threshold s_{thres}, we
predicted the consumption of food resource α by species i (or production of metabolite α by
species i) to be true only if $s_{i\alpha} > s_{\text{thres}}$ (or $s_{\alpha i} > s_{\text{thres}}$). We achieved excellent performance in
inferring either food-microbe consumption interactions (green line and dots with AUROC=0.9
in Fig. 5d) or microbe-metabolite production interactions (red line and dots with AUROC=0.92
in Fig. 5d).

We then performed the same inference on real data from the avocado study[28]. The
results are shown in Fig. 5e. (Inference results of other studies are provided in the
Supplementary Tables.) Our results shown in Fig. 5e are in agreement with prior biological
knowledge that *Faecalibacterium prausnitzii* is a stronger producer of butyrate[47] than
Ruminococcus callidus, and *R. calidus* is a stronger producer of acetate than *F. prausnitzii*[48,49].

The inference results also enable us to construct the tripartite food-microbe-metabolite
graph. For the sake of simplicity, here we visualize the avocado-microbe-butyrate subgraph
(Fig. 5f). Note that increased butyrate levels have been shown to be beneficial to host health
by enhancing immune status[19–21]. For the avocado-microbe-butyrate subgraph, we focused on
the top-20 avacado-microbe consumption and top-20 microbe-butyrate production interactions
ranked by their absolute sensitivity values. Only nodes and links associated with these
interactions were shown in this subgraph. Widths of individual edges in this figure are
proportional to the absolute values of the corresponding sensitivities and node sizes for
microbes are proportional to the products of edge widths connecting this microbe to avocado
at the top and butyrate at the bottom of this subgraph. We ordered microbial nodes in the
middle layer in the increasing order of node sizes from left to right (Fig. 5f). This organization
helps us identify the key species that serve as both strong consumers of avocado and strong
producers of butyrate. *F. prausnitzii* emerged as the most important key species for butyrate

production in response to avocado intervention. Our results are consistent with previous
studies[47]. For example, *F. prausnitzii* levels have been previously shown to be elevated when
avocado is supplied by diet[50]. In a separate study, *F. prausnitzii* has also been shown to
produce butyrate as a metabolic byproduct[47].
303

Discussion

A highly accurate computational method for predicting metabolic responses based on baseline
data and a potential dietary intervention strategy is a prerequisite for precision nutrition. In this
paper, we developed a deep learning method, McMLP, which predicts metabolomic profiles
after a dietary intervention better than existing methods. We first validated the superior
performance of McMLP using synthetic data generated by a microbial consumer-resource
model and investigated the influence of diet intervention doses and training sample sizes. We
then demonstrated that McMLP produced the most accurate predictions across six different
dietary intervention studies[27,28,35–38]. We proceeded with a biological interpretation of McMLP
results using sensitivity analysis to infer the tripartite food-microbe-metabolite relationship,
finding that the inferred relationship was quite accurate in synthetic data. Finally, we
demonstrated that our sensitivity analysis applied to real data revealed key species whose
metabolic capabilities were consistent with prior biological knowledge.

Currently available dietary intervention studies have many limitations for use in
machine learning. First, the sample size (or number of participants) of these studies is typically
small, on the order of dozens[27,32,37,38]. The relatively small sample size fundamentally limits the
performance of any predictive model. This problem may be mitigated in ongoing large-scale
research cohorts with many participants. One such cohort is the All of Us Research Program,
which is attempting to build a diverse health database of more than one million people across
the U.S. and then use the data to learn how our biology, lifestyle, and environment affect health.
As part of this observational cohort, the recently announced Nutrition for Precision Health Study
will recruit 10,000 participants to conduct precision dietary interventions[51]. Second, only a
handful of dietary components have ever been the subject of a dedicated diet-microbiota
studies. As a result, the computational approaches can only predict metabolic responses for
the limited set of dietary components used in these studies. However, to realize the promise of
precision nutrition to provide accurate personalized dietary recommendations, we need a
predictive model that can accurately predict metabolic responses for a wide range of dietary
components. Last, other baseline variables unavailable to us here (e.g., meal composition, age,

sex, demographics, and anthropometric data) might help to improve the predictive
performance. If such data are available, they can be incorporated into McMLP as extra input
variables.

Our McMLP architecture is quite generic --- its input variables and their dimensions can
be easily adapted to fit more complex datasets. For example, if a particular dietary intervention
study documents an extensive list of dietary components, McMLP can be modified to include
an input node for each dietary component to reflect the amount and frequency of its
consumption. Similarly, the predicted output variables of McMLP need not be limited to
metabolomic profiles measured in fecal samples. It can be generalized to predict other
variables such as immune biomarkers or metabolite concentrations from blood samples.

Unlike other machine learning methods that typically require hyperparameter tuning to
achieve the best performance for each dataset with a different set of hyperparameters, McMLP
consistently outperformed existing machine learning methods across six real datasets even
without hyperparameter tuning. We speculate that McMLP exploited the recently observed
"double-descent" behavior for the risk curve[52], which suggests that an overparametrized deep-
learning model (i.e., one with an extremely large number of model parameters) can generate
better and more consistent performance than models with less capacity and more carefully
tuned hyperparameters. To reach this overparameterized regime, we used a large and fixed
number of layers $N_l = 6$ and a large hidden layer dimension $N_h = 2048$, exceeding both the
number of microbial species and the number of metabolites. One benefit of using such a model
free of hyperparameter tuning is the shorter training time. Since the typical 5-fold cross-
validation used to select the best set of hyperparameters is the most time-consuming part of a
typical deep learning workflow, McMLP saves a significant amount of time required for
hyperparameter tuning and thus has a shorter training time (~ 5 minutes for each run of McMLP
on the avocado intervention study[28]).

357
358

Methods

Datasets.

The datasets utilized herein were generated as part of work on bacterial[53] and metabolite[54]
biomarkers of food intake, which provided anonymized microbial and metabolomic data on
Github. The main characteristics of the dietary intervention studies used above are
summarized in Table 1. Across all studies, fecal or blood samples were collected before and

after each dietary intervention period. Gut microbiota composition was determined by the 16S

rRNA gene sequencing and metabolomic profiles of either fecal samples or blood serum

samples were determined by tandem liquid chromatography-mass spectrometry (LC-MS/MS)

and gas chromatography-mass spectrometry (GC-MS) metabolomics. For all machine learning

tasks, the same five random 80/20 train-test splits were used to ensure a fair comparison of

methods. Further details are described below:

Avocado intervention study. This dataset was reported by a dietary intervention study

that investigated how avocado consumption altered the relative abundance of gut bacteria and

concentrations of microbial metabolites in 132 overweight or obese adults[28]. All participants

were assigned to the avocado treatment or no-avocado control group (66 each for arm). They

consumed isocaloric meals with or without avocado (175 g, men; 140 g, women) once daily for

12 weeks. For fecal samples collected before and after the dietary intervention, 278 ASVs

(Amplicon Sequence Variants) were determined by the 16S rRNA gene sequencing and

profiles of 6 SCFAs and 21 bile acids were generated by LC-MS/MS metabolomics.

Grains intervention study. This dietary intervention study investigated how grain barley

and oat consumption affects gut bacteria relative abundances and concentrations of microbial

metabolites in 68 healthy adults[35]. All participants were randomly assigned to receive one of

three treatments: (1) a control diet containing 0.8 daily servings of whole grain/1800 kcal, (2)

a diet containing 4.4 daily servings of whole grain barley/1800 kcal or (3) a diet containing 4.4

daily servings of whole grain oats/1800 kcal. Fecal samples were collected before and after

the dietary intervention.

Walnut intervention study. This dietary intervention study investigated how walnut

consumption affects the gut microbiota and metabolite concentrations in 18 healthy adults[27].

All participants completed two 3-week treatment/intervention periods separated by a 1-week

washout period. Fecal samples were collected before and after the dietary intervention period.

Almond intervention study. This dietary intervention study was conducted in 18 healthy

adults[36]. All participants completed four 3-week treatment periods and one control period

separated by a 1-week washout period. Fecal samples were collected before and after the

dietary intervention period.

Broccoli intervention study. In this study, 18 healthy adults completed two 18-day

treatment periods separated by a 24-day washout period[37]. Fecal samples were collected

before and after the dietary intervention period.

Fibers or fermented foods intervention study. This dietary intervention study was

designed to investigate how consumption of plant-based foods rich in dietary fibers or

fermented foods alters gut bacteria and their associated metabolites in 36 healthy adults[38]. All

participants were divided to the high-fiber or the high-fermented-foods arm (18 each for arm).

The entire dietary intervention lasted 17 weeks. Their fecal or blood serum samples were

collected before and after the dietary intervention period. Gut microbiota composition in fecal

samples was determined by the 16S rRNA gene sequencing and metabolomic profiles of

serum samples were generated by the LC-MS metabolomics.

McMLP. McMLP consists of two coupled MLPs: (step-1) in the first step (using the MLP at the

top in Supplementary Fig. 1a), we predict endpoint microbial compositions based on baseline

microbial compositions, baseline metabolomic profiles, and dietary intervention strategy; (step-

2) in the second step (using the MLP at the bottom in Supplementary Fig. 1a), we take the

predicted endpoint microbial compositions from the first MLP, baseline metabolomic profiles,

and dietary intervention strategy to predict endpoint metabolomic profiles.

•   Data processing: The CLR (Centered Log-Ratio) transformation is applied to microbial

relative abundances and the log10 transformation is applied to metabolite

concentrations.

•   Model detail: Each MLP model (for either the top or the bottom MLP in Supplementary

Fig. 1) has 6 hidden layers in the middle, sandwiched by input and output variables.

Each hidden layer has a fixed hidden layer dimension of 2048.

•   Training method: The Adam optimizer[55] is used for the gradient descent. Training stops

when the mean SCC (Spearman Correlation Coefficient) of annotated metabolites $\bar{\rho}$ on

the training set is less than 0.1 and $\bar{\rho}$ on the validation/test set starts to decrease within

the last 20 epochs.

•   Activation function: ReLU (Rectified Linear Unit).

Inference of food-microbe and microbe-metabolite interactions via sensitivity. The

two MLP models in the well-trained McMLP can be interpreted separately. We first interpret

the first MLP (step 1) in McMLP for food-microbe consumption interactions by the amount of

food resource α (Δx_α) and then measure the change in the relative abundance of species

i (Δy_i). Mathematically, for the sample m in the training set, we set the new value of this

variable as zero. As a result, the perturbation amount for this variable in sample m is $\Delta x_\alpha^{(m)} =$

$0 - x_\alpha^{(m)} = -x_\alpha^{(m)}$ where $x_\alpha^{(m)}$ is the unperturbed value. We can measure the change in the

relative abundance of species i for sample m ($\Delta y_i^{(m)}$) and define the sensitivity of species i to

food resource α for sample m as $s_{i\alpha}^{(m)} = \frac{\Delta y_i^{(m)}}{\Delta x_\alpha^{(m)}}$. Finally, we can average sensitivity values

across samples to obtain the average sensitivity of species i to food resource α: $s_{i\alpha} = \frac{\sum_m s_{i\alpha}^{(m)}}{N_{\text{train}}}$

where N_{train} is the number of training samples. Similarly, for the second MLP (step-2) in

McMLP, we can define $s_{\beta i}^{(m)} = \frac{\Delta y_\beta^{(m)}}{\Delta x_i^{(m)}}$ and $s_{\beta i} = \frac{\sum_m s_{\beta i}^{(m)}}{N_{\text{train}}}$ to infer microbe-metabolite interactions

by perturbing the relative abundance of species i (Δx_i) and then measuring the change in

concentration of metabolite β (Δy_β).

Statistics. To calculate correlations throughout the study, we used Spearman's correlation

coefficient. Wherever P-values were used we calculated the associated null distributions were

computed from scratch. All statistical tests were performed using standard numerical and

scientific computing libraries in the Python programming language (version 3.7.1) and Jupyter

Notebook (version 6.1).

Data and code availability. All code for the simulations used in this manuscript can be found

at https://github.com/wt1005203/McMLP.

Acknowledgements. Y.-Y.L. acknowledges grants from the National Institutes of Health

(R01AI141529, R01HD093761, RF1AG067744, UH3OD023268, U19AI095219, and

U01HL089856).

Author contributions. Y.-Y.L. conceived the project. T.W. and Y.-Y.L. designed the project.

T.W. performed all the numerical calculations and data analysis. All authors analyzed the

results. T.W. and Y.Y.L. wrote the manuscript. All authors edited and approved the manuscript.

Correspondence. Correspondence and requests for materials should be addressed to Y.-Y.

L. (email: yyl@channing.harvard.edu).

Competing Interests. The authors declare no competing interests.

References

1.   Hurtado-Lorenzo, A., Honig, G. & Heller, C. Precision Nutrition Initiative: Toward Personalized Diet Recommendations for Patients With Inflammatory Bowel Diseases. *Crohn's & Colitis 360* **2**, otaa087 (2020).

2.   Bush, C. L. *et al.* Toward the Definition of Personalized Nutrition: A Proposal by The American Nutrition Association. *Journal of the American College of Nutrition* **39**, 5–15 (2020).

3.   Zeevi, D. *et al.* Personalized Nutrition by Prediction of Glycemic Responses. *Cell* **163**, 1079–1094 (2015).

4.   Berry, S. E. *et al.* Human postprandial responses to food and potential for precision nutrition. *Nat Med* **26**, 964–973 (2020).

5.   Asnicar, F. *et al.* Microbiome connections with host metabolism and habitual diet from 1,098 deeply phenotyped individuals. *Nat Med* **27**, 321–332 (2021).

6.   Kolodziejczyk, A. A., Zheng, D. & Elinav, E. Diet–microbiota interactions and personalized nutrition. *Nat Rev Microbiol* **17**, 742–753 (2019).

7.   Chen, L. *et al.* Influence of the microbiome, diet and genetics on inter-individual variation in the human plasma metabolome. *Nat Med* 1–11 (2022) doi:10.1038/s41591-022-02014-8.

8.   Holscher, H. D. Dietary fiber and prebiotics and the gastrointestinal microbiota. *Gut Microbes* **8**, 172–184 (2017).

9.   Donia, M. S. & Fischbach, M. A. Small molecules from the human microbiota. *Science* **349**, 1254766 (2015).

10.   Koh, A., De Vadder, F., Kovatcheva-Datchary, P. & Bäckhed, F. From Dietary Fiber to Host Physiology: Short-Chain Fatty Acids as Key Bacterial Metabolites. *Cell* **165**, 1332–1345 (2016).

11.   Koppel, N., Maini Rekdal, V. & Balskus, E. P. Chemical transformation of xenobiotics by the human gut microbiota. *Science* **356**, eaag2770 (2017).

12.   Myhrstad, M. C. W., Tunsjø, H., Charnock, C. & Telle-Hansen, V. H. Dietary Fiber, Gut Microbiota, and Metabolic Regulation—Current Status in Human Randomized Trials. *Nutrients* **12**, 859 (2020).

13.   Cong, J., Zhou, P. & Zhang, R. Intestinal Microbiota-Derived Short Chain Fatty Acids in Host Health and Disease. *Nutrients* **14**, 1977 (2022).

14.   Corrêa-Oliveira, R., Fachi, J. L., Vieira, A., Sato, F. T. & Vinolo, M. A. R. Regulation of immune cell function by short-chain fatty acids. *Clin Transl Immunology* **5**, e73 (2016).

15.   Parada Venegas, D. *et al.* Short Chain Fatty Acids (SCFAs)-Mediated Gut Epithelial and Immune Regulation and Its Relevance for Inflammatory Bowel Diseases. *Frontiers in Immunology* **10**, (2019).

16.   Dalile, B., Van Oudenhove, L., Vervliet, B. & Verbeke, K. The role of short-chain fatty acids in microbiota–gut–brain communication. *Nat Rev Gastroenterol Hepatol* **16**, 461–478 (2019).

17.   Richards, L. B., Li, M., van Esch, B. C. A. M., Garssen, J. & Folkerts, G. The effects of short-chain fatty acids on the cardiovascular system. *PharmaNutrition* **4**, 68–111 (2016).

18.   Ohira, H., Tsutsui, W. & Fujioka, Y. Are Short Chain Fatty Acids in Gut Microbiota Defensive Players for Inflammation and Atherosclerosis? *Journal of Atherosclerosis and Thrombosis* **24**, 660–672 (2017).

19.   Segain, J.-P. *et al.* Butyrate inhibits inflammatory responses through NFκB inhibition:
implications for Crohn's disease. *Gut* **47**, 397–403 (2000).
20.   Säemann, M. D. *et al.* Anti-inflammatory effects of sodium butyrate on human
monocytes: potent inhibition of IL-12 and up-regulation of IL-10 production. *The FASEB*
Journal **14**, 2380–2382 (2000).
21.   He, J. *et al.* Intestinal butyrate-metabolizing species contribute to autoantibody
production and bone erosion in rheumatoid arthritis. *Science Advances* **8**, eabm1511 (2022).
22.   Hughes, R. L., Marco, M. L., Hughes, J. P., Keim, N. L. & Kable, M. E. The Role of
the Gut Microbiome in Predicting Response to Diet and the Development of Precision
Nutrition Models—Part I: Overview of Current Methods. *Advances in Nutrition* **10**, 953–978
(2019).
23.   Hughes, R. L., Kable, M. E., Marco, M. & Keim, N. L. The Role of the Gut
Microbiome in Predicting Response to Diet and the Development of Precision Nutrition
Models. Part II: Results. *Advances in Nutrition* **10**, 979–998 (2019).
24.   Senior, M. Precision nutrition to boost cancer treatments. *Nature Biotechnology* **40**,
1422–1424 (2022).
25.   David, L. A. *et al.* Diet rapidly and reproducibly alters the human gut microbiome.
Nature **505**, 559–563 (2014).
26.   Riaz Rajoka, M. S. *et al.* Interaction between diet composition and gut microbiota and
its impact on gastrointestinal tract health. *Food Science and Human Wellness* **6**, 121–130
(2017).
27.   Holscher, H. D. *et al.* Walnut Consumption Alters the Gastrointestinal Microbiota,
Microbially Derived Secondary Bile Acids, and Health Markers in Healthy Adults: A
Randomized Controlled Trial. *The Journal of Nutrition* **148**, 861–867 (2018).
28.   Thompson, S. V. *et al.* Avocado Consumption Alters Gastrointestinal Bacteria
Abundance and Microbial Metabolite Concentrations among Adults with Overweight or
Obesity: A Randomized Controlled Trial. *The Journal of Nutrition* **151**, 753–762 (2021).
29.   Suez, J. & Elinav, E. The path towards microbiome-based metabolite treatment. *Nat*
Microbiol **2**, 1–5 (2017).
30.   Wong, A. C. & Levy, M. New Approaches to Microbiome-Based Therapies. *mSystems*
4, e00122-19 (2019).
31.   Gulliver, E. L. *et al.* Review article: the future of microbiome-based therapeutics.
Alimentary Pharmacology & Therapeutics **56**, 192–208 (2022).
32.   Deehan, E. C. *et al.* Precision Microbiome Modulation with Discrete Dietary Fiber
Structures Directs Short-Chain Fatty Acid Production. *Cell Host & Microbe* **27**, 389-404.e6
(2020).
33.   Wang, D. D. & Hu, F. B. Precision nutrition for prevention and management of type 2
diabetes. *The Lancet Diabetes & Endocrinology* **6**, 416–426 (2018).
34.   Wastyk, H. C. *et al.* Gut-microbiota-targeted diets modulate human immune status.
Cell **184**, 4137-4153.e14 (2021).
35.   Thompson, S. V., Swanson, K. S., Novotny, J. A., Baer, D. J. & Holscher, H. D.
Gastrointestinal Microbial Changes Following Whole Grain Barley and Oat Consumption in
Healthy Men and Women. *The FASEB Journal* **30**, 406.1-406.1 (2016).
36.   Holscher, H. D., Taylor, A. M., Swanson, K. S., Novotny, J. A. & Baer, D. J. Almond
Consumption and Processing Affects the Composition of the Gastrointestinal Microbiota of
Healthy Adult Men and Women: A Randomized Controlled Trial. *Nutrients* **10**, 126 (2018).

37.   Kaczmarek, J. L. *et al.* Broccoli consumption affects the human gastrointestinal
microbiota. *The Journal of Nutritional Biochemistry* **63**, 27–34 (2019).
38.   Wastyk, H. C. *et al.* Gut-microbiota-targeted diets modulate human immune status.
Cell **0**, (2021).
39.   Mallick, H. *et al.* Predictive metabolomic profiling of microbial communities using
amplicon or metagenomic sequences. *Nature Communications* **10**, 1–11 (2019).
40.   Reiman, D., Layden, B. T. & Dai, Y. MiMeNet: Exploring microbiome-metabolome
relationships using neural networks. *PLOS Computational Biology* **17**, e1009021 (2021).
41.   Le, V., Quinn, T. P., Tran, T. & Venkatesh, S. Deep in the Bowel: Highly
Interpretable Neural Encoder-Decoder Networks Predict Gut Metabolites from Gut
Microbiome. *BMC Genomics* **21**, 256 (2020).
42.   Wang, T. *et al.* Predicting metabolomic profiles from microbial composition through
neural ordinary differential equations. *Nature Machine Intelligence* (in press).
43.   Zierer, J. *et al.* The fecal metabolome as a functional readout of the gut microbiome.
Nat Genet **50**, 790–795 (2018).
44.   Belkin, M., Hsu, D., Ma, S. & Mandal, S. Reconciling modern machine-learning
practice and the classical bias–variance trade-off. *PNAS* **116**, 15849–15854 (2019).
45.   Marsland, R. *et al.* Available energy fluxes drive a transition in the diversity, stability,
and functional structure of microbial communities. *PLoS Comput Biol* **15**, e1006793 (2019).
46.   An Introduction to Statistical Learning | SpringerLink.
https://link.springer.com/book/10.1007/978-1-0716-1418-1.
47.   Zhou, L. *et al.* Faecalibacterium prausnitzii Produces Butyrate to Maintain Th17/Treg
Balance and to Ameliorate Colorectal Colitis by Inhibiting Histone Deacetylase 1.
Inflammatory Bowel Diseases **24**, 1926–1940 (2018).
48.   Cremon, C. *et al.* Effect of Lactobacillus paracasei CNCM I-1572 on symptoms, gut
microbiota, short chain fatty acids, and immune activation in patients with irritable bowel
syndrome: A pilot randomized clinical trial. *United European Gastroenterology Journal* **6**,
604–613 (2018).
49.   Gaffney, J., Embree, J., Gilmore, S. & Embree, M. 2022. Ruminococcus bovis sp.
nov., a novel species of amylolytic Ruminococcus isolated from the rumen of a dairy cow.
International Journal of Systematic and Evolutionary Microbiology **71**, 004924.
50.   Cires, M. J. *et al.* Protective Effect of an Avocado Peel Polyphenolic Extract Rich in
Proanthocyanidins on the Alterations of Colonic Homeostasis Induced by a High-Protein
Diet. *J. Agric. Food Chem.* **67**, 11616–11626 (2019).
51.   NIH awards $170 million for precision nutrition study. *National Institutes of Health*
(NIH) https://www.nih.gov/news-events/news-releases/nih-awards-170-million-precision-
nutrition-study (2022).
52.   Belkin, M., Hsu, D., Ma, S. & Mandal, S. Reconciling modern machine learning
practice and the bias-variance trade-off. *arXiv:1812.11118 [cs, stat]* (2019).
53.   Shinn, L. M. *et al.* Fecal Bacteria as Biomarkers for Predicting Food Intake in Healthy
Adults. *The Journal of Nutrition* **151**, 423–433 (2021).
54.   Shinn, L. M. *et al.* Fecal Metabolites as Biomarkers for Predicting Food Intake by
Healthy Adults. *The Journal of Nutrition* nxac195 (2022) doi:10.1093/jn/nxac195.
55.   Kingma, D. P. & Ba, J. Adam: A Method for Stochastic Optimization. Preprint at
https://doi.org/10.48550/arXiv.1412.6980 (2017).
593

Figure 1: **A typical dietary intervention study design.** Before the dietary intervention, the
baseline gut microbial compositions and metabolomic profiles (of either fecal samples or blood
samples) are measured. During the dietary intervention, one or a few dietary resources are
introduced (represented here by avocado) in addition to the baseline diet. The task we intend
to solve is to predict personalized metabolic responses after dietary intervention based on the
baseline gut microbial compositions, baseline metabolomic profiles, and the dietary
intervention strategy.

Figure 2: **The workflow of McMLP.** We aim to predict endpoint metabolomic profiles (i.e.,
metabolomic profiles after the dietary interventions) based on the baseline microbial
compositions (i.e., microbial compositions before the dietary intervention), dietary intervention
strategy, and baseline metabolomic profiles. Here we used a hypothetical example with n=5
training samples and 2 samples in the test set. For each sample, we considered N_s microbial
species, N_d dietary resources, and N_m metabolites. Across three panels, microbial species
and their relative abundances are colored blue, dietary resources and their intervention doses
are colored green, and metabolites and their concentrations are colored red. Icons associated
with baseline/endpoint data are bounded by solid black/dashed lines respectively. **a,** The
model architecture of McMLP. McMLP comprises two coupled MLPs. The first MLP at the top
(step 1) predicts the endpoint microbial compositions based on the baseline data and the
dietary intervention strategy. The predicted endpoint microbial compositions from the first MLP
are then provided as input to the second MLP at the bottom (step 2). The second MLP
combines the predicted endpoint microbial compositions, the dietary intervention strategy, and
the baseline metabolomic profiles to finally predict the endpoint metabolomic profiles. Details
of both MLPs can be found in Supplementary Fig. 1 and Methods. **b,** McMLP takes two types
of baseline data (baseline microbial compositions and baseline metabolomic profiles) and the
dietary intervention strategy as input variables and is trained to predict corresponding endpoint
metabolomic profiles. During training, the endpoint microbial composition is needed to train the
first MLP. By contrast, the second MLP directly takes the predicted endpoint microbial
composition instead of the actual endpoint microbial composition. **c,** The well-trained McMLP
can generate predictions for metabolomic profiles for the test set. During testing, no endpoint
microbial composition is needed because the second MLP directly takes the predicted
endpoint microbial composition from the first MLP as the input.

Figure 3: **McMLP provides better predictive power than previously developed computational methods for predicting endpoint metabolomic profiles on synthetic data generated from microbial consumer-resource models.** Three computational methods are compared: Random Forest (RF), Gradient Boosting Regressor (GBR), and McMLP. For each method, we either included ("w/ b" label) or did not include ("w/o b" label) baseline metabolomic profiles as input variables. Each method with a particular combination of input data is colored the same way in all panels. Standard errors are computed based on five random train-test splits and shown in all panels (as solid black vertical lines or transparent areas around their means). To compare different methods, we adopted three metrics: the mean Spearman Correlation Coefficient (SCC) $\bar{\rho}$, the fraction of metabolites with SCCs greater than 0.5 (denoted as $f_{\rho>0.5}$), and the mean SCC of the top-5 predicted metabolites $\bar{\rho}_5$. Error bars denote the standard error (n=5). **a1-a3,** For the synthetic data with intervention dose of 3 and 50 training samples, McMLP provides the best performance for all three metrics regardless of whether the baseline metabolomic profiles are included or not. **b1-b3,** When the intervention dose is 3, the predictive performance of all methods gets better and closer to each other as the training sample size increases. Including baseline metabolomic profiles also helps to improve the prediction. **c1-c3,** When 200 training samples are used, the performance gap between including and not including baseline metabolomic profiles shrinks as the intervention dose increases.

Figure 4: **McMLP is superior to previous methods in terms of predicting endpoint metabolomic profiles on real data from six dietary intervention studies.** Three computational methods are compared: Random Forest (RF), Gradient Boosting Regressor (GBR), and McMLP. For each method, we either included ("w/ b" label) or did not include ("w/o b" label) baseline metabolomic profiles as input variables. Each method with a particular combination of input data is colored the same in all panels. Standard errors are computed based on five random train-test splits and shown in all panels (solid black vertical lines). To compare different methods, we adopted three metrics: the mean Spearman Correlation Coefficient (SCC) $\bar{\rho}$, the fraction of metabolites with SCCs greater than 0.5 (denoted as $f_{\rho > 0.5}$), and the mean SCC of the top-5 predicted metabolites $\bar{\rho}_5$. Error bars denote the standard error (n=5). **a1-a3,** Comparison of the performance in predicting SCFAs on the data from the avocado intervention study[28]. **b1-b3,** Comparison of performance in predicting bile acids on the data from the avocado intervention study[28]. **c1-c3,** Comparison of predictive performance on the data from the grain intervention study[35]. **d1-d3,** Comparison of predictive performance on the data from the walnut intervention study[27]. **e1-e3,** Comparison of predictive performance on the data from the almond intervention study[36]. **f1-f3,** Comparison of predictive performance on the data from the broccoli intervention study[37]. **g1-g3,** Comparison of predictive performance on the data from the high-fiber food or fermented food intervention study[38].

Figure 5: **Applying sensitivity analysis of McMLP accurately infers food-microbe consumption interactions and microbe-metabolite production interactions in both synthetic and real data. a,** The sensitivity of the relative abundance of species i to the supplied dietary resource α is denoted as $s_{i\alpha}$. It is defined as the ratio between the change in the relative abundance of species i (Δy_i) and a small perturbation in the supplied dietary resource α (Δx_α). Similarly, the sensitivity of the concentration of metabolite β to the relative abundance of species i is denoted as $s_{\beta i}$. It is defined as the ratio between the change in the concentration of metabolite β (Δy_β) and the perturbation in the relative abundance of species i (Δx_i). **b,** The sensitivity values for food-microbe consumption interactions (colored in green) and microbe-metabolite production interactions (colored in red) in the synthetic data. **c,** The ground-truth food-microbe consumption rates (colored in green) and microbe-metabolite production rates (colored in red) in the synthetic data. **d,** The Area Under the Receiver Operating Characteristic (AUROC) curve based on True Positive (TP) rates and False Positive (FP) rates which are obtained by using different sensitivity thresholds to classify interactions. **e,** The sensitivity values for avocado-microbe consumption interactions (colored in green) and

microbe-metabolite production interactions (colored in red) for the real data from the avocado intervention study. **f,** The avocado-microbe-butyrate tripartite graph constructed based on the sensitivity values of avocado-microbe consumption interactions and microbe-butyrate production interactions for the real data from the avocado intervention study. The edge width and edge arrow sizes are proportional to the absolute values of the sensitivities. All microbes in the middle layer are arranged from left to right in the increasing order of the incoming edge width multiplied by the outgoing edge width.

Dietary Intervention Studies	# of participants	# of intervention periods/groups	# of ASVs	# of metabolites
Avocado[28]	132	2	278	27
Grains[35]	68	3	650	43
Walnut[27]	18	2	419	41
Almond[36]	18	5	714	43
Broccoli[37]	18	2	855	35
Fibers or fermented foods[38]	32	2	503	9

692 Table 1: **Summary of key features of dietary intervention studies used in our method**
693 **comparison.** ASVs: Amplicon Sequence Variants.