

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A reference map of potential determinants for the human serum metabolome

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1962804> since 2024-03-20T13:39:47Z

Published version:

DOI:10.1038/s41586-020-2896-2

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

A reference map of potential determinants for the human serum metabolome

Noam Bar ^{# 1 2}, Tal Korem ^{# 1 2 3 4 5}, Omer Weissbrod ^{1 2 6}, David Zeevi ^{1 2 7}, Daphna Rothschild ^{1 2}, Sigal Leviatan ^{1 2}, Noa Kosower ^{1 2}, Maya Lotan-Pompan ^{1 2}, Adina Weinberger ^{1 2}, Caroline I Le Roy ⁸, Cristina Menni ⁸, Alessia Visconti ⁸, Mario Falchi ⁸, Tim D Spector ⁸; IMI DIRECT consortium; Jerzy Adamski ^{9 10 11}, Paul W Franks ^{12 13}, Oluf Pedersen ¹⁴, Eran Segal ^{15 16}

¹ Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel.

² Department of Molecular Cell Biology, Weizmann Institute of Science, Rehovot, Israel.

³ Department of Systems Biology, Columbia University Irving Medical Center, New York, NY, USA.

⁴ Department of Obstetrics and Gynecology, Columbia University Irving Medical Center, New York, NY, USA.

⁵ CIFAR Azrieli Global Scholars Program, CIFAR, Toronto, Ontario, Canada.

⁶ Department of Epidemiology, Harvard T.H. Chan School of Public Health, Boston, MA, USA.

⁷ Center for Studies in Physics and Biology, The Rockefeller University, New York, NY, USA.

⁸ Department for Twin Research & Genetic Epidemiology, King's College London, London, UK.

⁹ Research Unit Molecular Endocrinology and Metabolism, Genome Analysis Center, Helmholtz Zentrum München, German Research Center for Environmental Health, Neuherberg, Germany.

¹⁰ Lehrstuhl für Experimentelle Genetik, Technische Universität München, Freising-Weihenstephan, Germany.

¹¹ Department of Biochemistry, Yong Loo Lin School of Medicine, National University of Singapore, Singapore, Singapore.

¹² Lund University Diabetes Center, Department of Clinical Sciences, Lund University, Malmö, Sweden.

¹³ Harvard T.H. Chan School of Public Health, Boston, MA, USA.

¹⁴ The Novo Nordisk Foundation Center for Basic Metabolic Research, Faculty of Health and Medical Sciences, University of Copenhagen, Copenhagen, Denmark.

¹⁵ Department of Computer Science and Applied Mathematics, Weizmann Institute of Science, Rehovot, Israel. eran.segal@weizmann.ac.il.

¹⁶ Department of Molecular Cell Biology, Weizmann Institute of Science, Rehovot, Israel. eran.segal@weizmann.ac.il.

Contributed equally.

Abstract

The serum metabolome contains a plethora of biomarkers and causative agents of various diseases, some of which are endogenously produced and some that have been taken up from the environment¹. The origins of specific compounds are known, including metabolites that are highly heritable^{2,3}, or those that are influenced by the gut microbiome⁴, by lifestyle choices such as

smoking⁵, or by diet⁶. However, the key determinants of most metabolites are still poorly understood. Here we measured the levels of 1,251 metabolites in serum samples from a unique and deeply phenotyped healthy human cohort of 491 individuals. We applied machine-learning algorithms to predict metabolite levels in held-out individuals on the basis of host genetics, gut microbiome, clinical parameters, diet, lifestyle and anthropometric measurements, and obtained statistically significant predictions for more than 76% of the profiled metabolites. Diet and microbiome had the strongest predictive power, and each explained hundreds of metabolites—in some cases, explaining more than 50% of the observed variance. We further validated microbiome-related predictions by showing a high replication rate in two geographically independent cohorts^{7,8} that were not available to us when we trained the algorithms. We used feature attribution analysis⁹ to reveal specific dietary and bacterial interactions. We further demonstrate that some of these interactions might be causal, as some metabolites that we predicted to be positively associated with bread were found to increase after a randomized clinical trial of bread intervention. Overall, our results reveal potential determinants of more than 800 metabolites, paving the way towards a mechanistic understanding of alterations in metabolites under different conditions and to designing interventions for manipulating the levels of circulating metabolites.

Main

We used mass spectrometry to profile serum samples from 491 healthy individuals for whom we had previously collected extensive clinical, lifestyle, dietary, genetics and gut microbiome data¹⁰ (Methods, Extended Data Table 1). Our untargeted metabolomics analysis measured the levels of 1,251 metabolites, covering a wide range of biochemicals including lipids, amino acids, xenobiotics, carbohydrates, peptides and nucleotides, and approximately 30% unidentified compounds (Extended Data Fig. 1a, Methods, Supplementary Table 1). To classify unidentified metabolites and aid in biomarker discovery, we designed models that accurately predict the candidate biological pathway of the metabolites (Extended Data Fig. 2, Supplementary Note 1, Supplementary Table 2–5). Most metabolites we measured were prevalent across the cohort, including 498 metabolites that were detected in all samples and 1,104 metabolites that were detected in more than 50% of the samples (Extended Data Fig. 1b). After quality control (Methods), 475 individuals with high-quality data were included in the subsequent analyses.

To validate the accuracy of our metabolomic measurements, we compared the levels of creatinine and cholesterol to those obtained using standardized laboratory tests (Methods) that were performed independently on a second blood sample taken from the participants on the same visit, and found good agreement (creatinine, Pearson's $R = 0.87$; cholesterol, $R = 0.79$; Extended Data Fig. 1c, d). We further found that, for 20 participants for which samples were taken one week apart, the two profiles were significantly correlated (Spearman $\rho = 0.68 \pm 0.06$, median \pm s.d.); this was in contrast to samples from different participants, which showed no correlation (Spearman $\rho = 0.05 \pm 0.12$; Methods, Extended Data Fig. 1e). These results validate the reproducibility and accuracy of our data, are consistent with previous work showing long-term stability in the human metabolome¹¹, and confirm that this metabolic profile is a unique, person-specific signature.

Robust predictions of serum metabolites

We trained gradient-boosted decision trees¹² (GBDT) algorithms that predict metabolite levels in held-out individuals (Methods, Supplementary Note 2). GBDT systematically outperformed linear models (Lasso; Methods), with a median and maximum explained variance gain of 8.3% and 43.2%, respectively, for prediction with diet data, and 4.6% and 14.9%, respectively, for

microbiome data (Extended Data Fig. 3). Notably, our predictions for more than 76% of the metabolite groups tested were statistically significant with at least one feature group after multiple hypothesis correction (Methods). In total, 335 metabolites were significantly explained by diet-related features, and 182 by the microbiome (Fig. 1a, b). Our models explained more than 10% of the variance for 543 metabolite groups (median 10.2%; range 0–73.5%; Fig. 1d, Supplementary Table 6), and more than 50% of the variance for 17 metabolites.

We next checked, for each feature group, whether any type of metabolite was enriched with superior predictions (Fig. 1c, Methods). We found that clinical data better predicted metabolites that were classified as blood lipids, amino acids and peptides, as opposed to xenobiotics and unidentified compounds, on which it performed worse than on other metabolites. By contrast, microbiome data better explained levels of xenobiotics ($P < 10^{-4}$) and unidentified compounds ($P < 0.001$), highlighting its potential for explaining the origins of the large number of unidentified compounds. We further found that predictions based on clinical data were significantly correlated with those based on diet (Spearman's $\rho = 0.30$, $P < 10^{-20}$), and had a weaker correlation with predictions based on the microbiome ($R = 0.21$, $P < 10^{-11}$). Predictions based on microbiome data had the highest correlation with those based on diet ($R = 0.44$, $P < 10^{-20}$). Finally, we found that metabolites associated with genetics could not be predicted by other feature groups, and there was a weak correlation between the prediction accuracy of a model containing all other features ('full model', Methods) and the heritability of metabolites ($R = 0.09$, $P < 0.005$). Altogether, each feature group was particularly informative with respect to a different set of metabolites (Extended Data Figs. 4, 5a).

To estimate the relative predictive power of each feature group across all metabolites, we built models to predict the principal metabolomic components (Extended Data Fig. 5b). Diet had the largest predictive power, inferring 48.9% of the variance explained by a model containing all features (Methods), whereas lifestyle factors explained only 1.9% (Fig. 1e). Notably, the predictive power of microbiome data was 30.8% that of the full model. As a large portion of these predictions did not overlap with the predictions of other data, these results highlight the importance of microbiome data in predicting and potentially determining serum metabolite levels.

Replication in external cohorts

To test the robustness and reproducibility of the models based on gut microbiome data, we validated their accuracy in two geographically independent cohorts (Methods): 1,004 samples from healthy older participants from the UK (from the TwinsUK Registry⁷), and 245 samples from Northern European individuals with type 2 diabetes (from the IMI DIRECT cohort⁸; Extended Data Table 1). Validation data were not available to us while developing the prediction models, which were trained only on samples from the initial Israeli cohort. We obtained predictions for metabolites that had statistically significant predictions (false discovery rate (FDR) < 0.1) with $R^2 > 5\%$ in the Israeli cohort (107 metabolites in TwinsUK, 50 in IMI DIRECT), using only microbiome data from the validation cohorts. Notably, 95 out of 107 and 28 out of 50 predictions were replicated (FDR < 0.1) in the healthy TwinsUK cohort and in the IMI DIRECT cohort of patients with type 2 diabetes, respectively, including all top 60 predictions in the TwinsUK cohort (Fig. 2, Supplementary Tables 7, 8). We note that most of the replicated associations are accompanied by a reduction in effect size; this is expected, particularly as a result of study-specific biases. These results indicate that our models reveal robust associations between serum metabolites and the gut microbiome, despite differences between both the populations and the protocols and staff used to assemble these cohorts. Finally, most significant associations between metabolite levels and body mass index were also

replicated in the TwinsUK cohort¹³ with high accuracy ($R = 0.85$, $P < 10^{-10}$; Extended Data Fig. 5c, Supplementary Table 9).

Diet and microbiome models are independent

Because the diet modulates the gut microbiome¹⁴, we compared the explained variance of metabolites obtained by models based on either. Although some metabolites—mostly related to coffee consumption—were significantly predicted by both diet and microbiome data, many were not (Supplementary Table 10). Furthermore, adding microbiome data to a diet-based prediction model improved its accuracy in 66% of cases (median and maximum gain of 2.1% and 62.6% respectively; Supplementary Table 11), whereas adding permuted data reduced the performance in 82% of cases (median and maximum gain of -1.7% and 7.4% respectively; Extended Data Fig. 5d–f). Finally, 34 metabolites were significantly predicted only by microbiome data. Altogether, these results suggest that the gut microbiome may be modulating the production of many circulating metabolites, independent of diet.

We next used feature attribution analysis (SHAP⁹; Methods) to interpret these models, infer the drivers of each prediction, and examine interactions between different predictive factors (Extended Data Fig. 6, Supplementary Note 3). We found dozens of diet and bacterial features that were strongly predictive of blood metabolites in our models (Fig. 3a, Extended Data Fig. 7). Notably, the reported consumption of coffee (both long-term and short-term; Methods) was a stronger predictor than other dietary features for the levels of a large number of xenobiotics and unidentified compounds. These included metabolites from the xanthine metabolism pathway such as paraxanthine (diet prediction Pearson $R = 0.64$, $P < 10^{-20}$) and caffeine ($R = 0.68$, $P < 10^{-20}$), as previously reported¹⁵. These metabolites were also significantly predicted using microbiome data, with a Clostridiaceae species being the main predictor. Another strong feature was long-term fish consumption, which accurately predicted the levels of several blood lipids including 3-carboxy-4-methyl-5-propyl-2-furanpropionic acid (diet $R = 0.71$, $P < 10^{-20}$), a uraemic toxin that accumulates in the serum of patients with chronic kidney disease¹⁶ and has also been suggested to prevent and reverses steatosis¹⁷. X-16124 (microbiome $R = 0.77$, $P < 10^{-20}$) and X-11850 ($R = 0.7$, $P < 10^{-20}$) are two unidentified metabolites that were accurately predicted by microbiome data, and specifically by bacteria from the Eggerthellaceae family and Clostridium genus, respectively. Microbiome data was also highly predictive of the uraemic toxins phenylacetylglutamine ($R = 0.63$, $P < 10^{-20}$) and indoxyl sulfate ($R = 0.37$, $P < 10^{-20}$), which have been previously reported in association with cardiovascular disease¹⁸ and chronic kidney disease¹⁹; these predictions were driven by a species from the Lachnospiraceae family.

To assess whether a few important taxa are sufficient for accurate prediction, we defined the ‘main predictor’ of each metabolite as the taxon with the maximal mean absolute SHAP value. Nineteen bacterial taxa were the main predictors for the top 50 microbiome-predicted metabolites (prediction $R > 0.4$; Supplementary Table 12). One Clostridiaceae species was the main predictor of 22 of these, which are also strongly associated with coffee consumption in diet-based models. Clostridium sp. CAG:138 was the main predictor of five metabolites, including phenylacetylcarnitine ($R = 0.47$, $P < 10^{-20}$) and p-cresol-glucuronide ($R = 0.64$, $P < 10^{-20}$) as previously reported²⁰. Other taxa, however, were the main predictors of only one or two top metabolites, demonstrating that many different bacteria are required to accurately predict the levels of different metabolites. Among the main bacterial predictors of the top 100 metabolites, 89 belonged to Firmicutes, highlighting the strong predictive power of this phylum. It is notable that, although Bacteroidetes is the second most abundant phylum in our cohort (Extended Data Fig. 8a), none of its species were among these main predictors.

To check whether main predictors are sufficient for accurate prediction, for each metabolite we compared the accuracy of a full microbiome model to the accuracy of a model based only on its main predictor (Fig. 3b). We found that a model based on the main predictor could explain only a median of 36% of the explained variance of the full microbiome-based model. Cinnamoylglycine, for example, is significantly predicted using microbiome data ($R = 0.49$, $P < 10^{-20}$); however, a model based on its main predictor fails to provide a significant prediction. By contrast, some metabolites are exclusively predicted by a single bacterial species, such as the unidentified metabolite X-16124, for which a model based on an Eggerthellaceae species explained 93% of the variance of a full microbiome-based model. Indeed, in 95% of the individuals in which this bacteria was detectable, X-16124 was also detectable in serum, compared to only 23% of individuals for which this bacteria was not detected (Mann–Whitney U-test, $P < 10^{-20}$; Extended Data Fig. 8b).

New genetic–metabolomics associations

Several genome-wide association studies have found that human genetics influences serum metabolites^{2,3,21,22,23,24}. The median serum metabolite ACE-heritability, using the traditional twin model, was estimated to be 25%, whereas the median narrow-sense heritability, based only on discovered genetic loci, was estimated to be 2.1%. Because we measured several molecules that were not yet identified in these studies, we searched for associations between levels of these molecules and single nucleotide polymorphisms (SNPs; Supplementary Note 4). Notably, we found 68 statistically significant associations ($P < 5 \times 10^{-11}$ for all), of which—to the best of our knowledge (Methods)—22 have not been previously reported (Supplementary Table 13). These include ethylmalonate, a branched fatty acid that has been reported in association with anorexia nervosa²⁵ and that was associated with rs2066938, which explained 50% of its variance. This SNP is a variant of the 3'-untranslated region of the gene *UNC119B*, which we also found to be associated with butyrylcarnitine, in line with previous reports². Other examples include 2'-O-methyluridine and 2'-O-methylcytidine—both of which are nucleotides involved in pyrimidine metabolism—which we found to associate with a missense variant in the *PHYHD1* gene and have been previously reported to be negatively correlated with *PHYHD1* expression²⁶. We further found that X-21441—which we predicted as an androgenic steroid (Supplementary Note 1)—was associated with rs8187710, a missense variant in the *ABCC2* gene, explaining 11% of its variance (Extended Data Fig. 9). rs8187710 was previously demonstrated to be associated with non-alcoholic fatty liver disease²⁷. Notably, X-21441 was also negatively correlated with age in our cohort ($R = -0.3$, $P < 10^{-7}$), independent of the genotype (Extended Data Fig. 9c). This suggests that X-21441 might be an independent metabolic risk factor that mediates the genetic susceptibility to non-alcoholic fatty liver disease and age, a known risk factor for the disease²⁸.

Proof-of-concept clinical validation

As a proof-of-concept analysis, we examined whether some of the feature–metabolite interactions we uncovered may be causal. We used our diet-based models to select the top 5% of metabolites that were either positively or negatively associated with normal consumption of white or whole wheat bread (Fig. 4a, b, Methods). We then analysed the serum metabolome from the beginning and the end of a previously conducted week-long intervention²⁹, in which two randomized groups of ten healthy individuals increased their consumption of either wholegrain sourdough bread or industrial white bread (Fig. 4a, Methods). Notably, we found that metabolites that were positively associated with the consumption of whole wheat bread in our discovery cohort increased significantly more after the sourdough bread intervention (median fold change 1.62) than metabolites that were negatively associated with it (median fold change 0.66; Mann–Whitney U-test, $P < 10^{-10}$; Fig. 4c).

Moreover, we found no statistically significant differences when comparing the mean fold change of these metabolites under the white bread intervention ($P > 0.1$; Fig. 4c).

Some metabolites for which levels increased after the sourdough bread intervention were previously linked to the consumption of wholegrain wheat flour. A notable example is betaine, an amino acid that has been shown to improve vascular risk factors³⁰ and is also highly abundant in wheat bran and germ³¹. We found that the mean fold change in betaine levels in the sourdough bread group was 6.16, as opposed to 0.82 in the white bread group (Mann–Whitney U-test, $P < 0.004$; Fig. 4d). Another example is cytosine, for which the mean fold change was far greater in the sourdough bread group (78.5) compared to the white bread group (0.53) ($P < 0.002$; Fig. 4e) To the best of our knowledge, unlike betaine, cytosine levels have not previously been linked to bread consumption.

A similar analysis of metabolites that were associated with white bread consumption in our cohort did not find significant changes in their levels after intervention, potentially due to high baseline white wheat consumption in the typical diet of the study population. Overall, these results suggest that some of the associations that we found here might be causal.

Discussion

Although our cohort is not the largest in which serum metabolomics were measured, it is—to our knowledge—the only one in which these measurements were coupled with such a diverse array of potential determinants. Still, it has several limitations. First, although drug intake has been shown to have a large effect on the serum metabolome profile³², our cohort was healthy and had limited drug intake. We are therefore likely to be underestimating the potential effect of drug intake on blood metabolites. Second, replication of results is still required for predictions by most factors other than the microbiome. Third, owing to the lack of reliable annotations, we have not associated metabolites with specific enzymes; this could be addressed in subsequent experimental studies by focusing on strongly predictive taxa. Finally, because this study is mainly based on observational data, interpretation of interactions should be made with caution, and the associations cannot be considered as causal.

Taken together, our results reveal a comprehensive list of potential determinants for circulating blood metabolites. Many of the associations and interactions detected here replicate previously reported findings, supporting the validity of our results. The majority of them, however, are new, making them a useful resource for future studies, either for improving molecular understanding of health and disease, or for forming the basis of interventional studies aimed at altering the levels of blood metabolites.

Methods

No statistical methods were used to predetermine sample size. The experiments were not randomized and the investigators were not blinded to allocation during experiments and outcome assessment.

Description of cohorts

We analysed banked samples from two previously collected cohorts^{10,29}, for a total of 491 Israeli individuals. Studies were approved by Tel Aviv Sourasky Medical Center Institutional Review Board (IRB), approval numbers TLV-0658-12, TLV-0050-13 and TLV-0522-10; Kfar Shaul Hospital IRB, approval number 0–73. All participants signed written informed consent forms. Full

study designs, including inclusion and exclusion criteria are described elsewhere^{10,29}. In brief, participants in both studies were healthy individuals aged between 18 and 70. The participants answered detailed medical, lifestyle and nutritional questionnaires, provided stool and serum samples for metagenomic sequencing and metabolomics, were genotyped, underwent a comprehensive blood test, and for a period of at least one week, recorded all of their daily activities and nutritional intake in real-time using their smartphones with a specialized app provided to them²⁹. Both blood and stool samples were not taken under strict fasting conditions. Sixteen samples of participants for which microbiome data was not available to us were excluded from all analyses. Meetings in which participants provided blood samples took place in two different centres, Weizmann (45% of participants) and Tel-Aviv (55% of participants). All meetings in Weizmann took place within the first half of the day, whereas most meetings in Tel-Aviv took place during the second half of the day (82% of the participants).

Feature groups

The ‘diet’ feature group includes both answers for a detailed food frequency questionnaire (FFQ) aimed at capturing long term dietary habits, and the daily mean consumption of different food types, computed over a week based on real-time logging. In both cases we kept only items that were reported to be consumed at least once by at least 1% of our participants, resulting in 670 different food types from logging, and 141 different items from the FFQ.

The ‘macronutrients’ feature group includes the daily mean consumption of macronutrients (lipids, proteins, carbohydrates), calories and water, calculated from real-time logging.

The ‘anthropometrics’ feature group includes weight, BMI, waist and hips circumference, and waist-to-hips ratio.

The ‘cardiometabolic’ feature group includes systolic and diastolic blood pressure, heart rate in beats per minute and a glycemic status as previously described³³.

The ‘drugs’ feature group includes 30 binary features representing the intake of 20 common medications as reported in questionnaires, in addition to 10 medication groups as previously described³³. We included only drugs reported to be used by at least 1% of our participants.

The ‘clinical data’ feature group includes the age and sex of the participants, and the following feature groups described above: anthropometrics, cardiometabolic and drugs.

The ‘lifestyle’ feature group includes smoking status (current, past), stress levels obtained from questionnaires, and the daily mean sleeping time, exercise time and midday sleep time based on real time logging.

The ‘time of day’ feature is a binary feature indicating whether the sample was taken during the first half of the day.

The ‘seasonal effects’ feature is the month in which the sample was taken. In some analyses we also grouped months by season (Winter: December–February; Spring: March–May; Summer: June–August; Autumn: September–November).

The ‘microbiome’ feature group includes bacterial relative abundance calculated both by considering coverage (see the following section ‘Microbiome preprocessing’), and by MetaPhlAn234, as well as the first 10 principal components computed over the log transformed relative abundance of a bacterial gene catalogue³⁵ as previously described^{33,36}. Preprocessing steps are described in the following section.

We further defined a full model that included all of the above.

Metabolomics profiling and preprocessing

Metabolite concentrations were measured in serum samples by Metabolon, by using an untargeted liquid chromatography coupled to mass spectrometry (LC–MS) platform as previously described^{2,37,38}. A total of 540 serum samples were profiled, 19 of which were control samples (technical replicate) pooled from several individuals. The other 521 serum samples belonged to 491 participants.

We removed from further analysis 27 metabolites with fewer than 10 measurements across our cohort, and 54 metabolites that we found to have significantly different distributions in samples collected in two different recruitment centres (Mann–Whitney U-test, $P < 0.05/1,251$; Bonferroni corrected; Supplementary Table 14). For the remaining 1,170 metabolites, we performed robust standardization (subtracting the median and dividing by the standard deviation) over the log (base 10)-transformed levels, followed by clipping outlier samples which were further than 5 standard deviations. We next used two separate normalization schemes: one for single metabolites, which we subsequently used in the feature attribution analysis, and the second for metabolite groups, which we used for global and enrichment analyses.

For single metabolites, we regressed metabolite levels against storage times (only for metabolites present in at least 50 samples), and finally, imputed missing values as the minimum value per metabolite. For the second scheme, metabolites were grouped by correlation with a Spearman's ρ threshold of 0.85. This is done in order to handle possible bias resulting from uncertainty of metabolite assignments and a high rate of highly correlated mass spectrometry peaks, and resulted in 1,067 metabolite groups, 982 of which are singletons. The value of the metabolite group was set to the mean. The category of each metabolite group was assigned based on majority vote, where unidentified compounds were excluded from the vote unless all metabolites in the group were unidentified.

Microbiome preprocessing

Sample collection, DNA extraction, and sequencing of the samples in this study was previously described^{10,29,33}. In brief, we used only samples that were collected using swabs, filtered metagenomic reads containing Illumina adapters, filtered low-quality reads and trimmed low-quality read edges. We detected host DNA by mapping with GEM39 to the human genome (hg19) with inclusive parameters, and removed human reads. We subsampled all samples to have 10 million reads.

Bacterial relative abundance estimation was performed by mapping bacterial reads to species-level genome bins (SGB) representative genomes⁴⁰ (Supplementary Table 15). We selected all SGB representatives from groups with at least five genomes, and for these representative genomes kept only unique regions as a reference dataset. Mapping was performed using Bowtie2⁴¹ and abundance was estimated by calculating the mean coverage of unique genomic regions across the 50% most densely covered areas as previously described^{36,42}. Feature names include the lowest taxonomy level identified.

Comparing metabolomics to laboratory tests

We compared the levels of both creatinine and cholesterol, which we previously obtained via standard laboratory tests¹⁰ with their metabolomic levels. Because the tests were performed by two different laboratories, we centred the tests by reducing from the value of each sample the mean of all tests taken in the laboratory in which it was performed. We then performed a standardization of

the resulting measurements. The metabolomic profiling and the laboratory tests were performed on two samples taken at the same blood draw.

Correlation of metabolic profiles within and between individuals

We compared the Spearman's correlations between standardized metabolomic profiles of the same participant taken one week apart ($n = 20$) to correlations between standardized metabolomic profiles of different individuals ($n = 475$). Each pair of samples taken from the same participant was run in the same metabolomic batch. In the group of different individuals, only pairs of individuals from the same batch were included (resulting in a total of 3,835 such pairs), and were further stratified by sex.

Predictive models of metabolite groups

We used gradient boosting decision trees from the LightGBM (v.2.1.2) package¹², in order to predict the levels of 1,067 metabolite groups on the basis of 7 feature groups in held-out individuals. In order to estimate the EV of each metabolite group we ran a fivefold cross-validation (CV) model using each feature group as input, and evaluated the results using the coefficient of determination (R^2). For all prediction results except those based on human genetics (Methods) we computed 95% confidence intervals and P values via 1,000 iterations of bootstrapping⁴³. In each bootstrap iteration, we performed a random fivefold cross validation, in which in each fold we randomly sampled (with replacement) a group of participants from the training set to have the same size as the current training set. We next used this set to train our model and evaluated the performance of the model on the set of participants in the remaining fold. Finally, we computed the coefficient of determination between the measured values of the metabolite and the concatenation of the CV's predicted values as obtained from the bootstrapping iteration. We applied the Fisher transformation to the estimations of the explained variance we got from bootstrapping in order to induce normality⁴⁴, and then computed a standard error, and estimated the P values via the normal cumulative distribution function using the Wald test⁴⁵, such that our null hypothesis is that the explained variance should distribute normally with zero mean. Confidence intervals were computed empirically from the bootstrapping results. We corrected P values of predictions for multiple hypotheses using the Benjamini–Hochberg procedure over all feature groups (10% FDR). In all CV and bootstrapping runs we used a fixed and predetermined set of hyperparameters: for the microbiome and diet feature groups: `learning_rate = 0.005`, `max_depth = default`, `feature_fraction = 0.2`, `num_leaves = default`, `min_data_in_leaf = 15`, `metric = L2`, `early_stopping_rounds = None`, `n_estimators = 2000`, `bagging_fraction = 0.8`, `bagging_freq = 1`; for other feature groups: `learning_rate = 0.01`, `max_depth = 5`, `feature_fraction = 0.8`, `num_leaves = 25`, `min_data_in_leaf = 15`, `metric = L2`, `early_stopping_rounds = None`, `n_estimators = 200`, `bagging_fraction = 0.9`, `bagging_freq = 5`.

Human genetics based prediction models

To obtain the predictions based on human genetics, we used a similar fivefold CV scheme, in which in every fold we calculated the associations between SNPs and metabolite levels within the training fold, and then trained a model on only the top 10 SNPs that reached genome-wide significance (Bonferroni-adjusted). For folds in which no SNP reached the significance level, we assigned every sample in the test fold with the mean metabolite level of the training fold. Owing to high complexity and running time issues, P values and confidence intervals were not computed based on bootstrapping; rather, we estimated the P values of the Pearson's correlation between the true and predicted metabolite levels. Metabolites for which the R^2 was negative were assigned a P value of 1.

Testing for SNP associations with metabolites

Genotype processing and imputation of 413 individuals were described previously³³. We performed genome-wide associations for single metabolites ($n = 1,170$) and calculated the P value and the estimated effect sizes using plink (v.1.07). When declaring a genome-wide significance for the SNP–metabolite associations we used a conservative Bonferroni adjustment procedure to control for the false discovery rate due to the large number of SNPs tested ($P < (5 \times 10^{-8})/1,170$). We performed all genome wide associations using imputed genotypes.

For named molecules, their chemical identification, super and sub pathways are presented as provided by Metabolon. For unidentified molecules, super and sub pathways are estimated on the basis of our biological pathway classifier. We did our best to scan the available literature for known associations between genetic loci and metabolites before reporting an association as novel. The main resources included the GWAS Catalog⁴⁶ and the GWAS server^{2,22}.

Pathway category enrichment analysis

For each pathway category we used a Mann–Whitney U-test comparing the prediction accuracy of metabolites from that category compared to prediction accuracy of metabolites from other categories. Direction of enrichment was determined by the sign of the Mann–Whitney U-test statistic. We considered only metabolite groups for which at least one feature group had a significant prediction (after correcting for multiple hypotheses), resulting in 819 metabolite groups.

Validation of metabolite predictions based on microbiome

We validated the robustness of the associations between the gut microbiome composition and the levels of circulating metabolites in two independent cohorts in which we had access to both metagenomics sequencing. Serum metabolomics in these cohorts were performed using the same Metabolon platform that we used for the discovery cohort. The first validation cohort included 1,004 samples of healthy participants from the TwinsUK cohort⁷, for which there was an average of 0.9 ± 1.3 years gap between the collection of faecal and blood samples. The second validation cohort included 245 samples of participants of European ancestry with type 2 diabetes (T2D) from the IMI DIRECT consortium⁸. Data from both these validation cohorts were not available to us while developing the prediction models. The metagenomics samples from both cohorts went through the exact same analysis pipeline as our discovery cohort to extract the bacterial features that our prediction models were based on. We then applied our models on these data to obtain the metabolite predictions for both cohorts. Only metabolites that were significantly predicted based only on microbiome data with $R^2 > 5\%$ ($FDR < 0.1$) in our discovery cohort were considered for further analysis (107 metabolites out of 678 in TwinsUK, 50 metabolites out of 261 in IMI DIRECT). Within every validation cohort, we performed robust standardization (subtracting the median and dividing by the standard deviation) over the log (base 10) transformed levels, followed by clipping outlier samples which were further than 5 standard deviations, and finally, imputed missing values as the minimum value per metabolite. The analysis of these geographically distinct cohorts holds multiple potential sources of noise, including different methods, centres and staff involved in assembling these cohorts, as well as different cohort demographics, clinical manifestations, different genetic background and dietary and lifestyle preferences. Therefore, we defined a successful replication as one that restores the original ranking of the participants as dictated by the true levels of the metabolite in hand. Hence, for every metabolite, in each validation cohort, we computed the Spearman's correlation between its true levels and its predicted levels. A replication was considered significant if the FDR adjusted P value of the Spearman's correlation was lower than 0.1 and the correlation coefficient was strictly positive.

Feature attribution analysis

To explain the output of our machine learning models and find specific associations between features and metabolite levels, we used SHAP (Shapley additive explanations)⁴⁷, a recently introduced framework for interpreting predictions, which assigns each feature an importance value for a particular prediction. In brief, for a specific prediction, the SHAP value of a feature is defined as the change in the expected value of the output of the model when this feature is observed versus when it is missing. It is computed using a sum that represents the effect of each feature being added to the model averaged over all possible orderings of features being introduced. Shapley-value-based analysis in gut microbiome data was recently demonstrated to be useful, as it enabled the estimation of complex contributions of gut microbiome taxa to functional shifts, while maintaining global community composition properties⁴⁸.

Individual SHAP values were computed for held-out individuals in fivefold CV using the module TreeExplainer (v.0.24.0)^{9,49}, based on models trained only on features from the respective feature group. Before training, we standardized the levels of target metabolites, so that SHAP values from different models would be comparable (they are measured in the same units as the target). In each CV fold we ran a random hyperparameter search consistent of 10 iterations using the module RandomizedSearchCV from sklearn (v.0.20.4), and chose the best model for predicting the held-out individuals and computing SHAP values. In all feature attribution analyses we used the ungrouped list of 1,170 metabolites.

For every feature, we computed the mean absolute SHAP value across all instances in a specific model, reflecting the mean effect of each feature on the predictions and serving as a feature importance measure. We further used these values to compute directional mean absolute SHAP values, by multiplying them with the sign of the Spearman correlation between the population feature and the target. Here, positive values indicate that higher feature values lead, on average, to higher predicted values, whereas negative values indicate that lower feature values lead, on average, to higher predicted values.

When performing feature attribution analysis with gut microbiome data as input, we included only the relative abundance of SGB representative genomes as features, taking only features which were present in over 5% of the samples, resulting with 753 bacterial taxa. When using diet as input, we considered only features that were present in at least 5% of the samples, resulting with 398 food types from logging and items from the FFQ.

Comparing gradient boosting decision trees with a linear model

We compared the EV of every single metabolite obtained for a GBDT and a Lasso regression model. The EV of all models were calculated in fivefold CV, where in each fold we ran a hyperparameter search consistent of 10 iterations as described above. We used LightGBM as the GBDT model, and Lasso regression (sklearn, v.0.20.4) as the linear model, since its regularization scheme is better suited for a large number of features, as in the case of diet and gut microbiome composition. Because GBDT handles missing values well, we first imputed all missing values as the median of each feature to assure a fair comparison. When applying the models on the microbiome data, we used log₁₀-transformed values.

Estimating relative predictive power of feature groups

To estimate the relative predictive power of different feature groups we first applied a principal component analysis over the metabolite groups data to get the first 400 PCs which constitute more than 99% of the total variance in the data (Extended Data Fig. 5b). We then used fivefold CV

prediction models as described above to predict the PCs on the basis of the different feature groups independently. As baseline, we used the full model, which consists of all features combined to predict the levels of the PCs, and estimated the overall fraction of variance explained by:

$$\frac{\sum_{i=1}^{nPC} EV_i \times PC_i}{\sum_{i=1}^{nPC} PC_i},$$

where EV_i is the fraction of EV that the model recovers for PC i . PC_i is the fraction of variance that PC i explains out of the overall variation in the data. nPC is the number of the first PCs, those which capture the most variation. For the features we have collected, we defined this sum obtained for the full model as the total explainable variance in circulating blood metabolites. Next, for every feature group we computed a similar expression and calculated the relative predictive power by dividing this expression by that of the full model. The estimates we present are for $nPC = 15$, as the overall EV of the full model that we estimated using the first 15 PCs constitutes over 97% of the overall EV of the full model based on all 400 PCs.

Biological sub pathway prediction

We used gradient boosting decision trees from the LightGBM (v.2.1.2) package¹², in order to build a multiclass classifier to predict the biological sub pathways of metabolites as annotated by Metabolon. When developing the classifier, we considered only named metabolites from biological sub pathways which include over 10 metabolites each in our data, resulting with 28 sub pathways covering a total of 572 named molecules (sub pathway size range 11–44). The rationale behind this is that we tried to find the balance between covering as many metabolites and types of metabolites possible while keeping the number of classes reasonable.

We trained our model in a leave-one-out CV scheme, in which in every training fold we used 20% of the training samples as internal validation to perform an early stopping of 50 rounds. We then obtained a soft max of size 28 per metabolite, representing the probabilities of every metabolite being labelled as one of the 28 sub pathways. For the prediction of the unidentified molecules, we used a model trained once using all 572 metabolites. The features used for the training of the model included the normalized levels of metabolites across our main discovery cohort, the mean raw count of the metabolite and the fraction of missing values across the discovery cohort. In addition, to capture the associations between metabolites and their predictive features, we included the directional mean absolute SHAP values for every pair of metabolite–feature computed from the ‘full model’ as described above. The final vectors of probabilities were determined as an ensemble of three models: the first, trained only on the SHAP values; the second, trained only on metabolite levels, means and fraction of missing values; and the third, trained on all combined. Finally, the mean of these three models was computed.

When evaluating the performance of our classifier on the named labelled molecules, we concatenated all vectors of probabilities resulting from the leave-one-out procedure. For every sub pathway we computed a classification report including the classification precision ($TP/(TP + FP)$), recall ($TP/(TP + FN)$) and f1-score ($2 \times (\text{precision} \times \text{recall})/(\text{precision} + \text{recall})$), to account for the imbalanced class sizes. The overall accuracy was computed as the fraction of metabolites with correctly assigned labels out of all metabolites from all sub pathways which were included in the training phase. In all runs we used a fixed and predetermined set of hyperparameters (objective = multiclass, num_leaves = 25, max_depth = 4, learning_rate = 0.005, bagging_fraction = 0.8, feature_fraction = 0.8, bagging_freq = 1, bagging_seed = 2018, class_weight = balanced,

n_estimators = 2000, early_stopping_rounds = 50). TP, true positive; FP, false positive; FN, false negative.

Characterization of unidentified metabolites by metabolon

Characterization of unidentified metabolites was done as previously described²¹. In brief, identification of tentative structural features for unidentified biochemicals incorporates a detailed analysis of mass spec data, that is, gathering information such as the accurate monoisotopic mass, the elution time and fragmentation pattern of the primary ion, and correlation to other molecules. The accurate monoisotopic mass is used to identify a probable structural formula for the unidentified biochemical, which is then used to search against chemical structure databases. When a candidate structure fits the accurate monoisotopic mass and fragmentation data, an authentic standard is commercially purchased or synthesized (when possible). Conformation of a proposed structure is based on a match to three primary criteria, including co-elution with the unidentified molecule of interest, and a high degree match to both the accurate monoisotopic mass and fragmentation pattern.

Interaction networks

We used a graphical layout in order to visualize the associations of features with the levels of metabolites. The nodes are either metabolites or features, and the edges are the directional mean absolute SHAP values computed from models trained only on features from the respective feature group as described above. All networks were constructed using Cytoscape⁵⁰. The threshold for presenting SHAP values as edges was determined as 0.12, keeping the network sparse enough for convenience of visualization.

Analysis of bread intervention

In order to find the associations between metabolite levels and the consumption of both types of bread in the study cohort we computed the directional mean absolute SHAP values of the reported consumption of both white and whole-wheat bread for all metabolites. The SHAP values were computed in cross validation from models based only on the reported consumption of each type of bread. We ranked the metabolites according to their directional mean absolute SHAP value for each type of bread and used the top 5% positively and negatively driven metabolites for further analysis. The prediction models were constructed using 458 samples of distinct individuals, a subset of our cohort from which we excluded all samples of individuals which participated in the intervention study.

For each metabolite in every individual, we computed the fold change of metabolite levels between the samples taken at the end of the first week of intervention and the start of that week. Before computing fold change we imputed missing values with the minimum per metabolite and standardized their log (base 10) transformed levels. Furthermore, for each intervention group, we computed the mean fold change of every metabolite based on the 10 samples from that group. We then compared the mean fold change of the top 5% positively and negatively associated metabolites mentioned above within each intervention group by performing a rank sum test (two-sided Mann–Whitney U-test) over the mean fold change.

For comparing the fold change of betaine and cytosine between the two intervention groups, we used a two-sided Mann–Whitney U-test.

Linear mixed models-based estimates of the explained variance of metabolites using gut microbiome

For the in-sample estimation of EV for metabolites based on gut microbiome we used a linear mixed model framework that we had recently developed³³. In brief, we used GCTA⁵¹, a tool used in statistical genetics for the estimating of SNP-based genetic kinship. Instead of a matrix of host SNPs, as is commonly used in GCTA, we used a kinship matrix computed over the presence-absence of microbial species which were also used as features in the out-of-sample prediction models. We added the storage time as a covariate to the model. P values were computed using RL-SKAT⁵².

Statistical analysis

For all statistical analysis and prediction models we used Python 2.7.8 with the following packages: pandas 0.23.4, numpy 1.14.2, scikit-learn 0.20.4, scipy 1.1.0, shap 0.24.0, LightGBM 2.1.2.

References

- 1 Psychogios, N. et al. The human serum metabolome. PLoS ONE 6, e16957 (2011).
- 2 Shin, S.-Y. et al. An atlas of genetic influences on human blood metabolites. Nat. Genet. 46, 543–550 (2014).
- 3 Long, T. et al. Whole-genome sequencing identifies common-to-rare variants associated with human blood metabolites. Nat. Genet. 49, 568–578 (2017).
- 4 Wikoff, W. R. et al. Metabolomics analysis reveals large effects of gut microflora on mammalian blood metabolites. Proc. Natl Acad. Sci. USA 106, 3698–3703 (2009).
- 5 Xu, T. et al. Effects of smoking and smoking cessation on human serum metabolite profile: results from the KORA cohort study. BMC Med. 11, 60 (2013).
- 6 Playdon, M. C. et al. Comparing metabolite profiles of habitual diet in serum and urine. Am. J. Clin. Nutr. 104, 776–789 (2016).
- 7 Moayyeri, A., Hammond, C. J., Hart, D. J. & Spector, T. D. The UK Adult Twin Registry (TwinsUK Resource). Twin Res. Hum. Genet. 16, 144–149 (2013).
- 8 Koivula, R. W. et al. Discovery of biomarkers for glycaemic deterioration before and after the onset of type 2 diabetes: rationale and design of the epidemiological studies within the IMI DIRECT Consortium. Diabetologia 57, 1132–1142 (2014).
- 9 Lundberg, S. M., Erion, G. G. & Lee, S.-I. Consistent individualized feature attribution for tree ensembles. Preprint at <https://arxiv.org/abs/1802.03888v3> (2018).
- 10 Zeevi, D. et al. Personalized nutrition by prediction of glycemic responses. Cell 163, 1079–1094 (2015).
- 11 Yousri, N. A. et al. Long term conservation of human metabolic phenotypes and link to heritability. Metabolomics 10, 1005–1017 (2014).
- 12 Ke, G. et al. LightGBM: a highly efficient gradient boosting decision tree. In Advances in Neural Information Processing Systems 30 (eds Guyon, I. et al.) (Neural Information Processing Systems Foundation, 2017).
- 13 Cirulli, E. T. et al. Profound perturbation of the metabolome in obesity is associated with health risk. Cell Metab. 29, 488–500.e2 (2019).
- 14 David, L. A. et al. Diet rapidly and reproducibly alters the human gut microbiome. Nature 505, 559–563 (2014).
- 15 Azam, S., Hadi, N., Khan, N. U. & Hadi, S. M. Antioxidant and prooxidant properties of caffeine, theobromine and xanthine. Med. Sci. Monit. 9, BR325–BR330 (2003).

- 16 Tsutsumi, Y. et al. Renal disposition of a furan dicarboxylic acid and other uremic toxins in the rat. *J. Pharmacol. Exp. Ther.* 303, 880–887 (2002).
- 17 Prentice, K. J. et al. CMPF, a metabolite formed upon prescription omega-3-acid ethyl ester supplementation, prevents and reverses steatosis. *EBioMedicine* 27, 200–213 (2018).
- 18 Nemet, I. et al. A cardiovascular disease-linked gut microbial metabolite acts via adrenergic receptors. *Cell* 180, 862–877.e22 (2020).
- 19 Hung, S.-C., Kuo, K.-L., Wu, C.-C. & Tarng, D.-C. Indoxyl sulfate: A novel cardiovascular risk factor in chronic kidney disease. *J. Am. Heart Assoc.* <https://doi.org/10.1161/JAHA.116.005022> (2017).
- 20 Evenepoel, P., Meijers, B. K. I., Bammens, B. R. M. & Verbeke, K. Uremic toxins originating from colonic microbial metabolism. *Kidney Int.* 76, S12–S19 (2009).
- 21 Yousri, N. A. et al. Whole-exome sequencing identifies common and rare variant metabolic QTLs in a Middle Eastern population. *Nat. Commun.* 9, 333 (2018).
- 22 Suhre, K. et al. Human metabolic individuality in biomedical and pharmaceutical research. *Nature* 477, 54–60 (2011).
- 23 Gieger, C. et al. Genetics meets metabolomics: a genome-wide association study of metabolite profiles in human serum. *PLoS Genet.* 4, e1000282 (2008).
- 24 Kettunen, J. et al. Genome-wide association study identifies multiple loci influencing human serum metabolite levels. *Nat. Genet.* 44, 269–276 (2012).
- 25 Capo-chichi, C. D. et al. Riboflavin and riboflavin-derived cofactors in adolescent girls with anorexia nervosa. *Am. J. Clin. Nutr.* 69, 672–678 (1999).
- 26 Darst, B. F., Lu, Q., Johnson, S. C. & Engelman, C. D. Integrated analysis of genomics, longitudinal metabolomics, and Alzheimer’s risk factors among 1,111 cohort participants. *Genet. Epidemiol.* 43, 657–674 (2019).
- 27 Sookoian, S., Castaño, G., Gianotti, T. F., Gemma, C. & Pirola, C. J. Polymorphisms of MRP2 (ABCC2) are associated with susceptibility to nonalcoholic fatty liver disease. *J. Nutr. Biochem.* 20, 765–770 (2009).
- 28 Hamaguchi, M. et al. Aging is a risk factor of nonalcoholic fatty liver disease in premenopausal women. *World J. Gastroenterol.* 18, 237–243 (2012).
- 29 Korem, T. et al. Bread affects clinical parameters and induces gut microbiome-associated personal glycemic responses. *Cell Metab.* 25, 1243–1253.e5 (2017).
- 30 Olthof, M. R., van Vliet, T., Boelsma, E. & Verhoef, P. Low dose betaine supplementation leads to immediate and long term lowering of plasma homocysteine in healthy men and women. *J. Nutr.* 133, 4135–4138 (2003).
- 31 Craig, S. A. S. Betaine in human nutrition. *Am. J. Clin. Nutr.* 80, 539–549 (2004).
- 32 Liu, J. et al. Integration of epidemiologic, pharmacologic, genetic and gut microbiome data in a drug-metabolite atlas. *Nat. Med.* 26, 110–117 (2020).
- 33 Rothschild, D. et al. Environment dominates over host genetics in shaping human gut microbiota. *Nature* 555, 210–215 (2018).
- 34 Segata, N. et al. Metagenomic microbial community profiling using unique clade-specific marker genes. *Nat. Methods* 9, 811–814 (2012).
- 35 Li, J. et al. An integrated catalog of reference genes in the human gut microbiome. *Nat. Biotechnol.* 32, 834–841 (2014).
- 36 Zeevi, D. et al. Structural variation in the gut microbiome associates with host health. *Nature* 568, 43–48 (2019).

- 37 Evans, A. M. et al. High resolution mass spectrometry improves data quantity and quality as compared to unit mass resolution mass spectrometry in high-throughput profiling metabolomics. *Metabolomics* 4, 1000132 (2014).
- 38 Zierer, J. et al. The fecal metabolome as a functional readout of the gut microbiome. *Nat. Genet.* 50, 790–795 (2018).
- 39 Marco-Sola, S., Sammeth, M., Guigó, R. & Ribeca, P. The GEM mapper: fast, accurate and versatile alignment by filtration. *Nat. Methods* 9, 1185–1188 (2012).
- 40 Pasoli, E. et al. Extensive unexplored human microbiome diversity revealed by over 150,000 genomes from metagenomes spanning age, geography, and lifestyle. *Cell* 176, 649–662.e20 (2019).
- 41 Langmead, B. & Salzberg, S. L. Fast gapped-read alignment with Bowtie 2. *Nat. Methods* 9, 357–359 (2012).
- 42 Korem, T. et al. Growth dynamics of gut microbiota in health and disease inferred from single metagenomic samples. *Science* 349, 1101–1106 (2015).
- 43 Efron, B. & Tibshirani, R. J. *An Introduction to the Bootstrap* (Chapman and Hall, 1994).
- 44 Fisher, R. A. Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population. *Biometrika* 10, 507 (1915).
- 45 Wald, A. Sequential tests of statistical hypotheses. *Ann. Math. Stat.* 16, 117–186 (1945).
- 46 Buniello, A. et al. The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019. *Nucleic Acids Res.* 47, D1005–D1012 (2019).
- 47 Lundberg, S. & Lee, S.-I. A unified approach to interpreting model predictions. Preprint at <https://arxiv.org/abs/1705.07874v2> (2017).
- 48 Manor, O. & Borenstein, E. Systematic characterization and analysis of the taxonomic drivers of functional shifts in the human microbiome. *Cell Host Microbe* 21, 254–267 (2017).
- 49 SHAP <https://github.com/slundberg/shap>
- 50 Shannon, P. et al. Cytoscape: a software environment for integrated models of biomolecular interaction networks. *Genome Res.* 13, 2498–2504 (2003).
- 51 Yang, J., Lee, S. H., Goddard, M. E. & Visscher, P. M. GCTA: a tool for genome-wide complex trait analysis. *Am. J. Hum. Genet.* 88, 76–82 (2011).
- 52 Schweiger, R. et al. RL-SKAT: an exact and efficient score test for heritability and set tests. *Genetics* 207, 1275–1283 (2017).

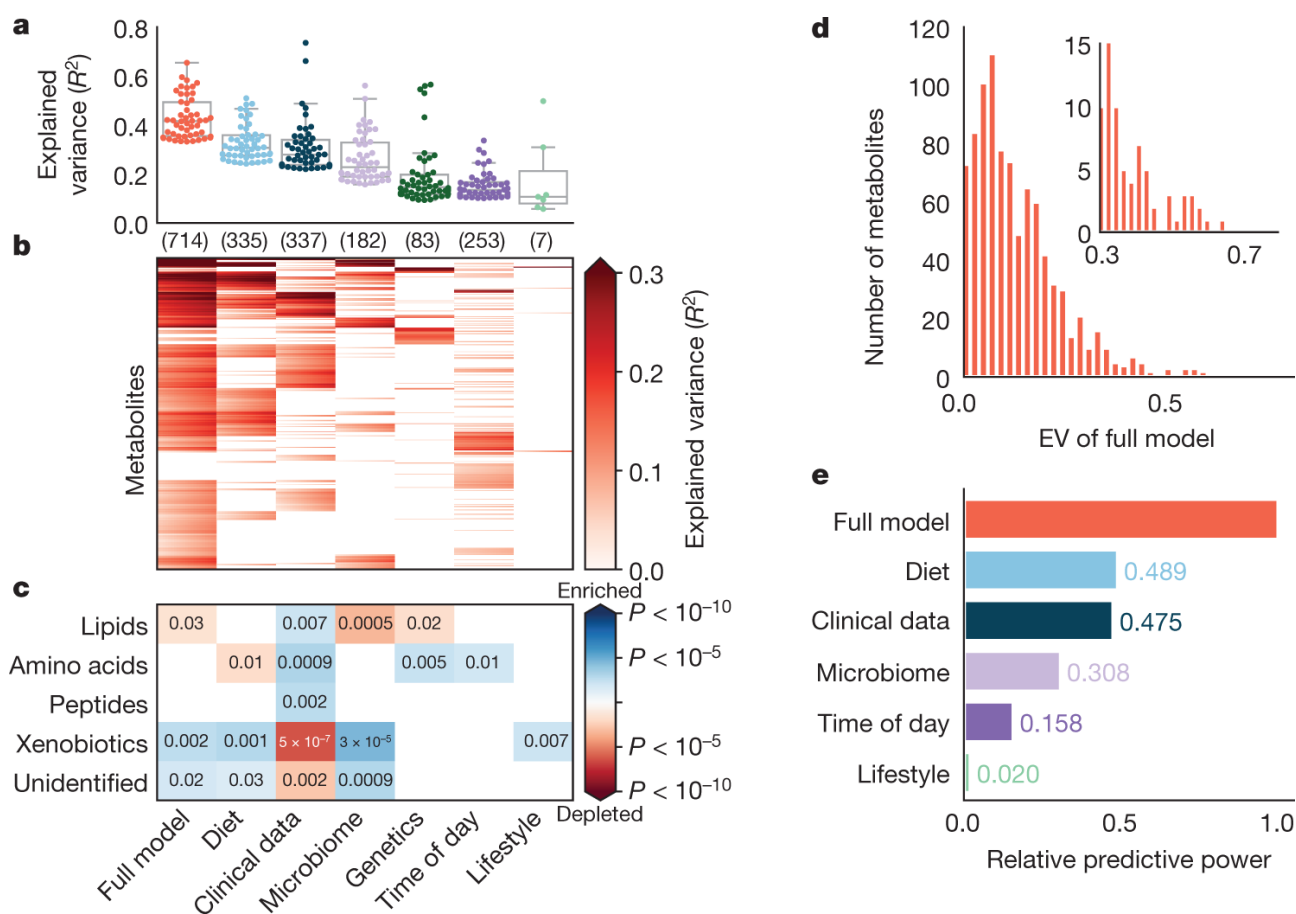
Consortia

The IMI DIRECT consortium

Henrik Vestergaard, Manimozhiyan Arumugam, Torben Hansen, Kristine Allin, Tue Hansen, Mun-Gwan Hong, Jochen Schwenk, Ragna Haussler, Matilda Dale, Toni Giorgino, Marianne Rodriguez, Mandy Perry, Rachel Nice, Timothy McDonald, Andrew Hattersley, Angus Jones, Ulrike Graefe-Mody, Patrick Baum, Rolf Grempler, Cecilia Engel Thomas, Federico De Masi, Caroline Anna Brorsson, Gianluca Mazzoni, Rosa Allesøe, Simon Rasmussen, Valborg Gudmundsdóttir, Agnes Martine Nielsen, Karina Banasik, Konstantinos Tsirigos, Birgitte Nilsson, Helle Pedersen, Søren Brunak, Tugce Karaderi, Agnete Troen Lundgaard, Joachim Johansen, Ramneek Gupta, Peter Wad Sackett, Joachim Tillner, Thorsten Lehr, Nina Scherer, Christiane Dings, Iryna Sihinevich, Heather Loftus, Louise Cabrelli, Donna McEvoy, Andrea Mari, Roberto Bizzotto, Andrea Tura, Leen 't Hart, Koen Dekkers, Nienke van Leeuwen, Roderick Sliker, Femke

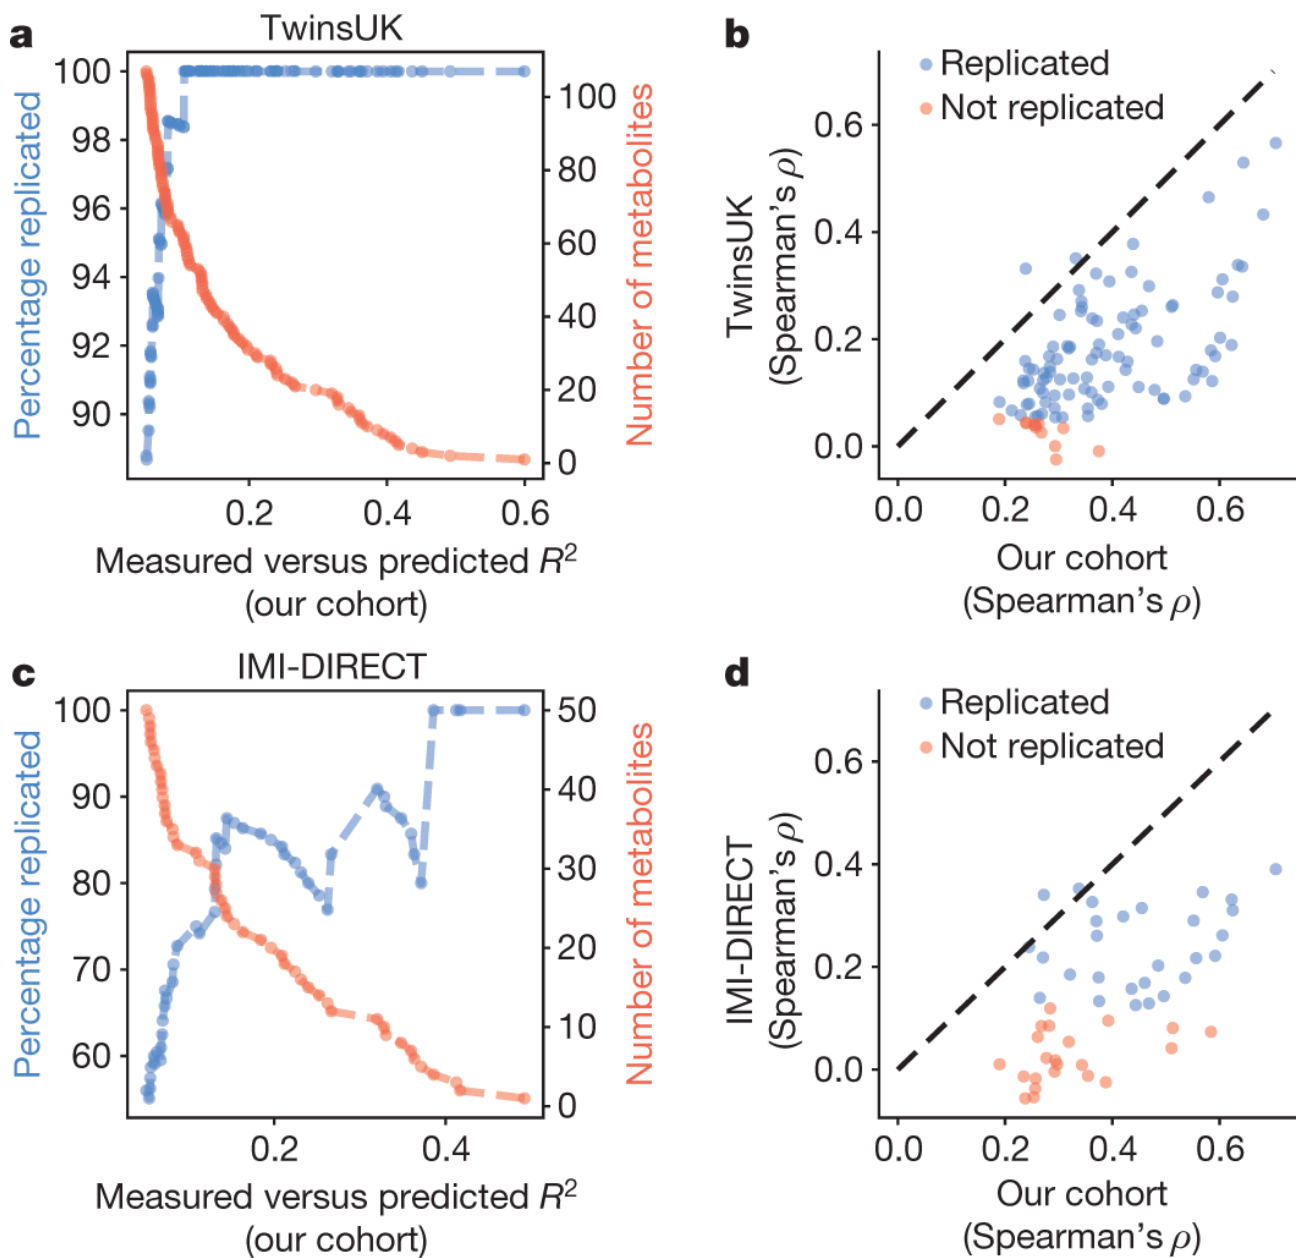
Rutters, Joline Beulens, Giel Nijpels, Anitra Koopman, Sabine van Oort, Lenka Groeneveld, Leif Groop, Petra Elders, Ana Viñuela, Anna Ramisch, Emmanouil Dermitzakis, Beate Ehrhardt, Christopher Jennison, Philippe Froguel, Mickaël Canouil, Amélie Boneford, Ian McVittie, Dianne Wake, Francesca Frau, Hans-Henrik Staerfeldt, Kofi Adragani, Melissa Thomas, Han Wu, Imre Pavo, Birgit Steckel-Hamann, Henrik Thomsen, Giuseppe Nicola Giordano, Hugo Fitipaldi, Martin Ridderstråle, Azra Kurbasic, Naeimeh Atabaki Pasdar, Hugo Pomares-Millan, Pascal Mutie, Robert Koivula, Nicky McRobert, Mark McCarthy, Agata Wesolowska-Andersen, Anubha Mahajan, Moustafa Abdalla, Juan Fernandez, Reinhard Holl, Alison Heggie, Harshal Deshmukh, Anita Hennige, Susanna Bianzano, Barbara Thorand, Sapna Sharma, Harald Grallert, Jonathan Adam, Martina Troll, Andreas Fritsche, Anita Hill, Claire Thorne, Michelle Hudson, Teemu Kuulasmaa, Jagadish Vangipurapu, Markku Laakso, Henna Cederberg, Tarja Kokkola, Yunlong Jiao, Stephen Gough, Neil Robertson, Helene Verkindt, Violeta Raverdi, Robert Caiazzo, Francois Pattou, Margaret White, Louise Donnelly, Andrew Brown, Colin Palmer, David Davtian, Adem Dawed, Ian Forgie, Ewan Pearson, Hartmut Ruetten, Petra Musholt, Jimmy Bell, Elizabeth Louise Thomas, Brandon Witcher, Mark Haid, Claudia Nicolay, Miranda Mourby, Jane Kaye, Nisha Shah, Harriet Teare, Gary Frost, Bernd Jablonka, Mathias Uhlen, Rebeca Eriksen, Josef Vogt, Avirup Dutta, Anna Jonsson, Line Engelbrechtsen, Annemette Forman, Nadja Sondertoft, Nathalie de Preville, Tania Baltauss, Mark Walker, Johann Gassenhuber, Maria Klintonberg, Margit Bergstrom & Jorge Ferrer

Fig. 1: Diet, gut microbiome, genetics and clinical data predict the levels of most serum metabolites.



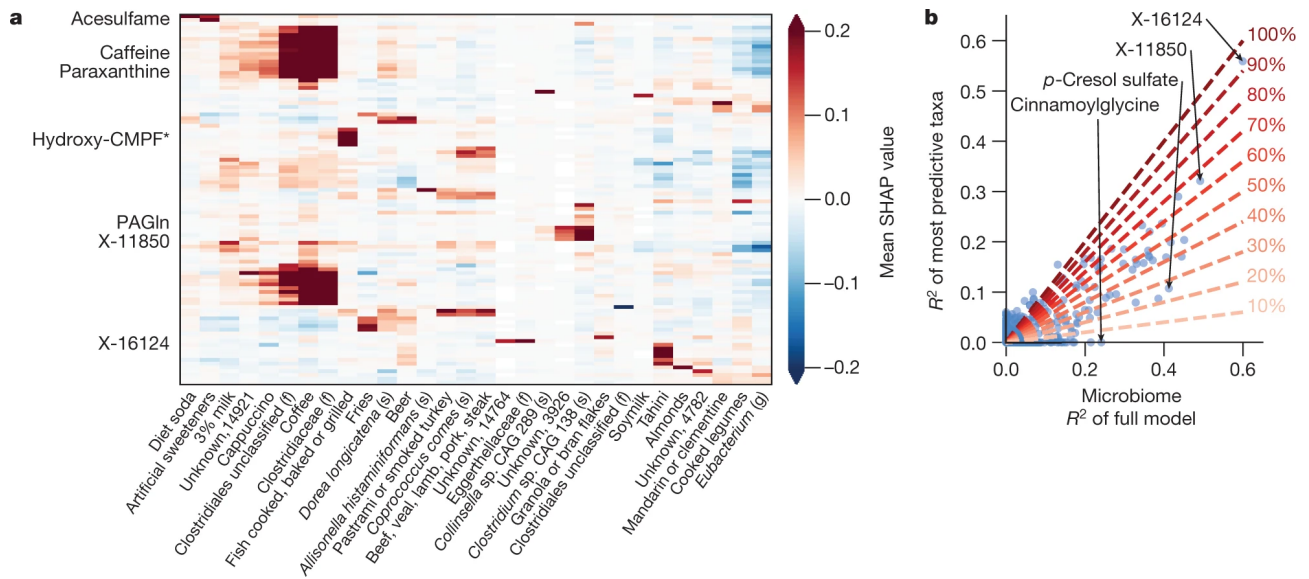
This dataset arises from fivefold cross-validation predictions of metabolite levels based on separate models for each feature group. **a**, Box and swarm plots (centre, median; box, interquartile range (IQR); whiskers, $1.5 \times \text{IQR}$) showing the explained variance (EV; R^2) of the top 50 significantly predicted metabolites (when available) of each feature group (group names are listed in **c**). **b**, Heat map with colour gradient from left to right corresponding to the 95% confidence interval for the EV, for each metabolite by every feature group. Only metabolites with significant predictions ($\text{FDR} < 10\%$) are shown, and the number of metabolites per group is shown in brackets at the top. P values and confidence intervals were estimated using bootstrapping (Methods). **c**, Enrichment of metabolite types in the predictions by each feature group (two-sided Mann–Whitney U-test; Methods). Only significant enrichments are shown ($P < 0.05$ after 10% FDR correction). Exact P values are written in each cell. **d**, Histogram of the number of metabolites with each value of EV obtained using the full model. The inset shows the EV range of 0.3–0.8. **e**, The fraction of total EV of each feature group compared to the total EV of a model with all feature groups excluding genetics (full model). The total EV is the sum of the EV of the first 15 metabolite principal components weighted by the EV of each principal component (Methods).

Fig. 2: Validation of microbiome-based predictions of metabolites in two independent cohorts.



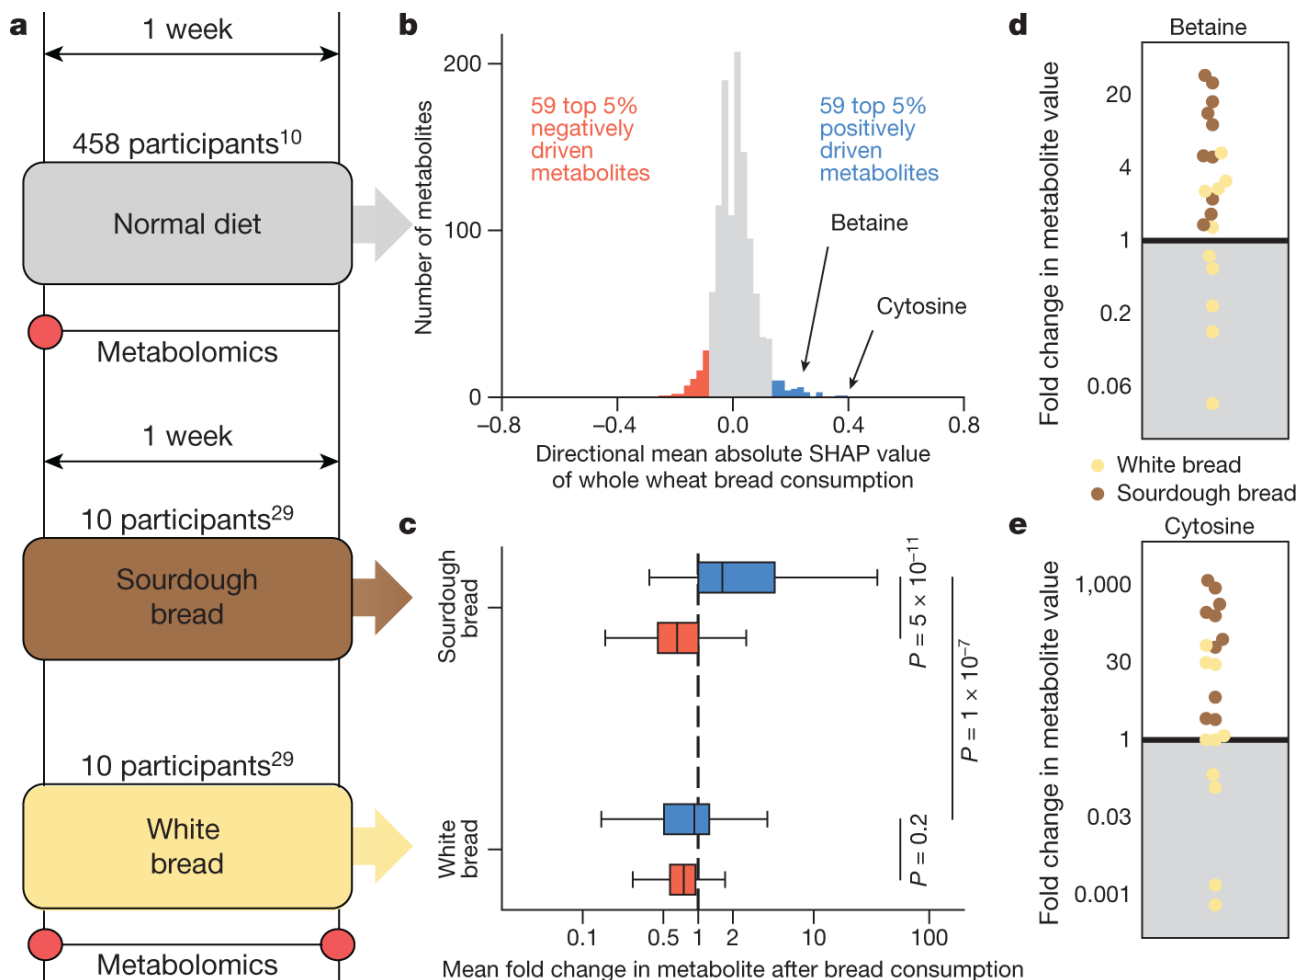
a, c, R^2 of predicted metabolites in our cohort plotted against the rate of replicated associations in the replication cohorts (TwinsUK, **a**; IMI DIRECT, **c**), computed as the fraction of significant replications out of all predictions with equal or higher predicted R^2 in our cohort (left y axis; blue; FDR < 10%), and the cumulative number of metabolites (right y axis; red). **b, d**, Spearman correlation between true and predicted levels of metabolites in our cohort plotted against the same correlations in the replication cohorts (TwinsUK, **b**; IMI DIRECT, **d**). Metabolites are coloured by the replication success (replicated, blue ($n = 95$ (**b**), $n = 28$ (**d**)); not replicated, red ($n = 12$ (**b**), $n = 22$ (**d**)); FDR < 10%).

Fig. 3: Diet and gut microbiome data independently explain a wide range of biochemicals.



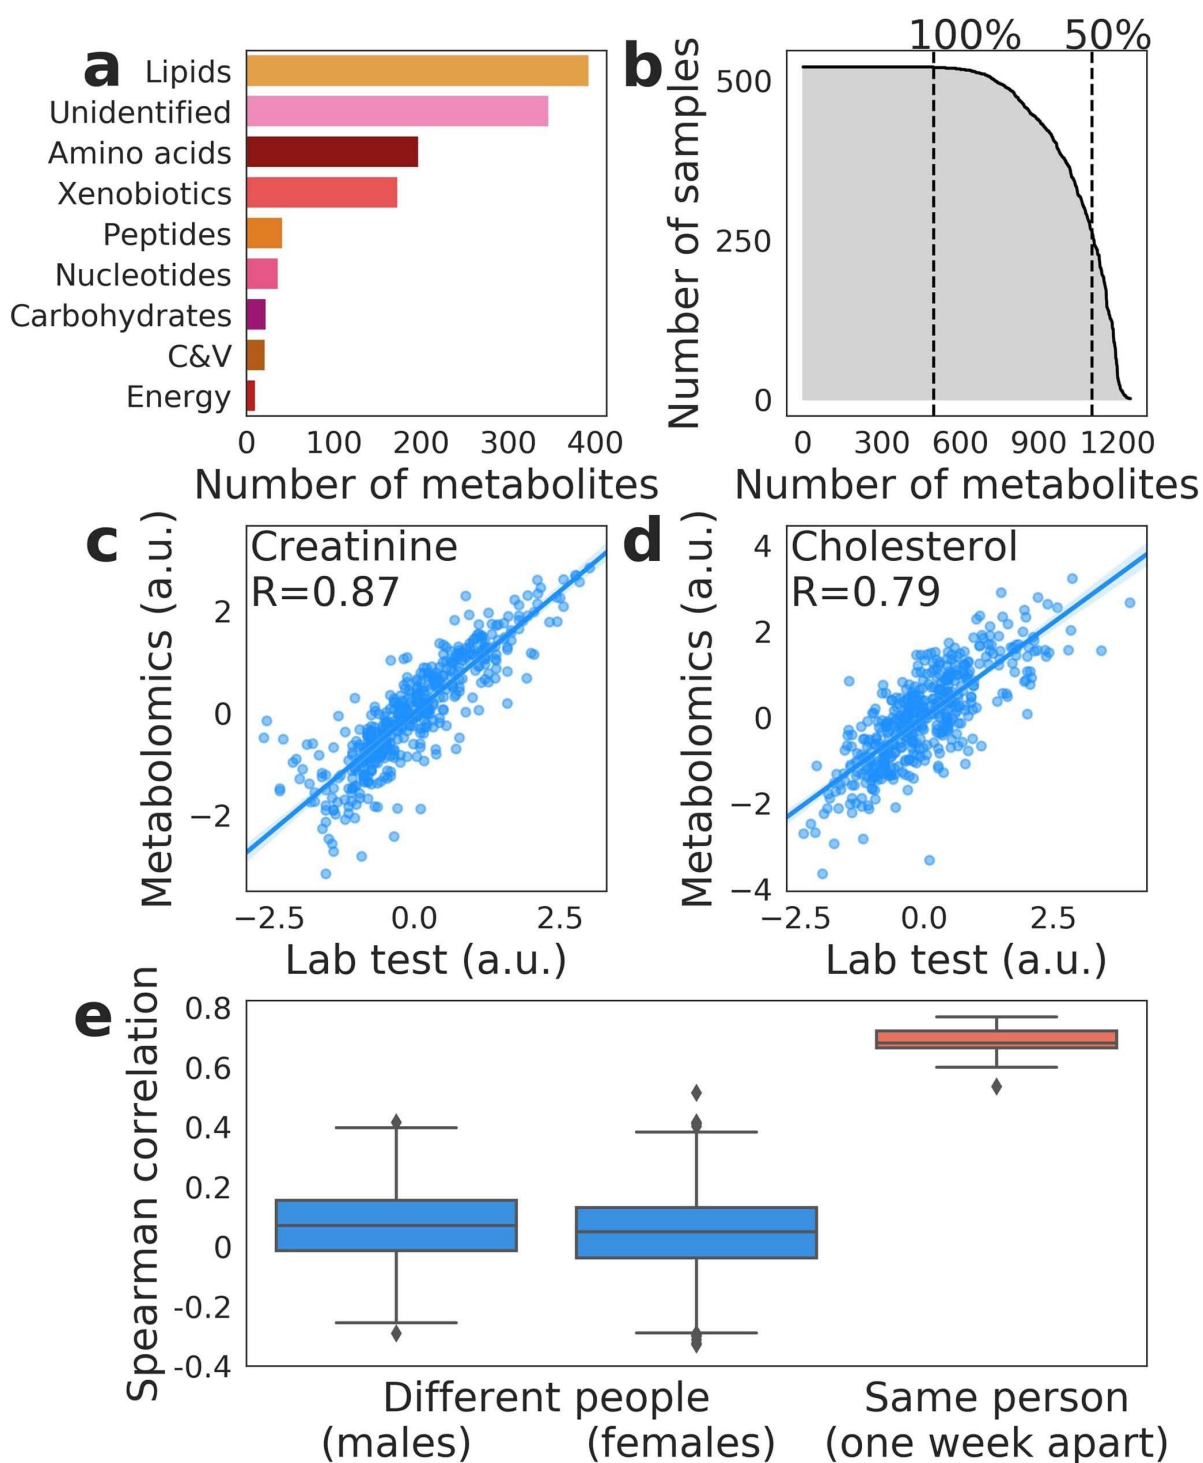
a, Subset of a heat map showing the directional mean absolute SHAP values (Methods) of various features (x axis) computed from fivefold cross-validation models that predict metabolite levels (y axis) using two separate models, one based on diet and another on gut microbiome data. Positive (negative) SHAP values indicate that higher (lower) feature values lead, on average, to higher predicted values. Shown are the top 100 predicted metabolites using diet and gut microbiome, and the top 30 features by maximum mean absolute SHAP value across all metabolites. See extended heat map in Extended Data Fig. 7g. **b**, The EV of every metabolite from microbiome-based prediction models (x axis) compared to using only the top predictor of that metabolite, selected as the feature with the largest mean absolute SHAP value (y axis). Dashed red lines mark different y:x ratios. PAGln, phenylacetylglutamine; hydroxy-CMPF*, hydroxy-3-carboxy-4-methyl-5-propyl-2-furanpropanoic acid.

Fig. 4: Increases in levels of metabolites explained by bread after an intervention of increased bread consumption.



a, Schematic of the measurement of metabolites and of white bread and whole wheat bread consumption. The prediction models were constructed using samples from distinct participants who consumed and recorded their normal diet for one week¹⁰ ($n = 458$ participants). We analysed samples from the first week of a randomized controlled trial²⁹, in which 10 participants increased their consumption of wholegrain sourdough bread and 10 others increased their consumption of industrial white bread. **b**, Histogram of directional mean absolute SHAP values of whole wheat bread consumption for metabolites computed on the basis of held-out samples from our cohort. The top 5% ($n = 59$; blue) positively associated metabolites and the top 5% ($n = 59$; red) negatively associated metabolites are used for further analysis. **c**, Box plots (centre, median; box, IQR; whiskers, $1.5 \times$ IQR) showing the mean fold change in the top 5% positively (blue) and negatively (red) associated metabolites, separated by intervention group. They show a significantly higher mean fold change for the top 5% positively associated metabolites compared with the top 5% negatively associated metabolites under the sourdough bread intervention (two-sided Mann–Whitney U-test, $P = 5 \times 10^{-11}$). **d**, **e**, The fold changes of metabolite levels for both betaine (**d**; two-sided Mann–Whitney U-test, $P = 0.0036$) and cytosine (**e**; $P = 0.0014$) were higher in the sourdough bread group compared to the white bread group.

Extended Data Fig. 1: Accurate and reproducible untargeted serum metabolomics.



a, Breakdown of the 1,251 measured metabolites by type. **b**, Number of samples (y axis) in which each metabolite (x axis) was identified, sorted by prevalence. **c**, **d**, Mass-spectrometry measurements (y axis) versus standardized lab tests results (x axis; Methods) for creatinine (**c**; Pearson's $R = 0.87$, $P < 10^{-20}$) and cholesterol (**d**; $R = 0.79$, $P < 10^{-20}$). **e**, Spearman correlations (y axis: centre, median; box, IQR; whiskers, $1.5 \times$ IQR) between standardized metabolomic profiles (Methods) of different individuals ($n = 475$; median Spearman's ρ 0.05, s.d. 0.12) stratified by sex, and between standardized metabolomic profiles of the same participant ($n = 20$; median Spearman ρ 0.68, s.d. 0.06) taken one week apart. C&V, cofactors and vitamins; a.u., arbitrary units.

Extended Data Fig. 2: Biological sub pathway prediction of unidentified molecules.

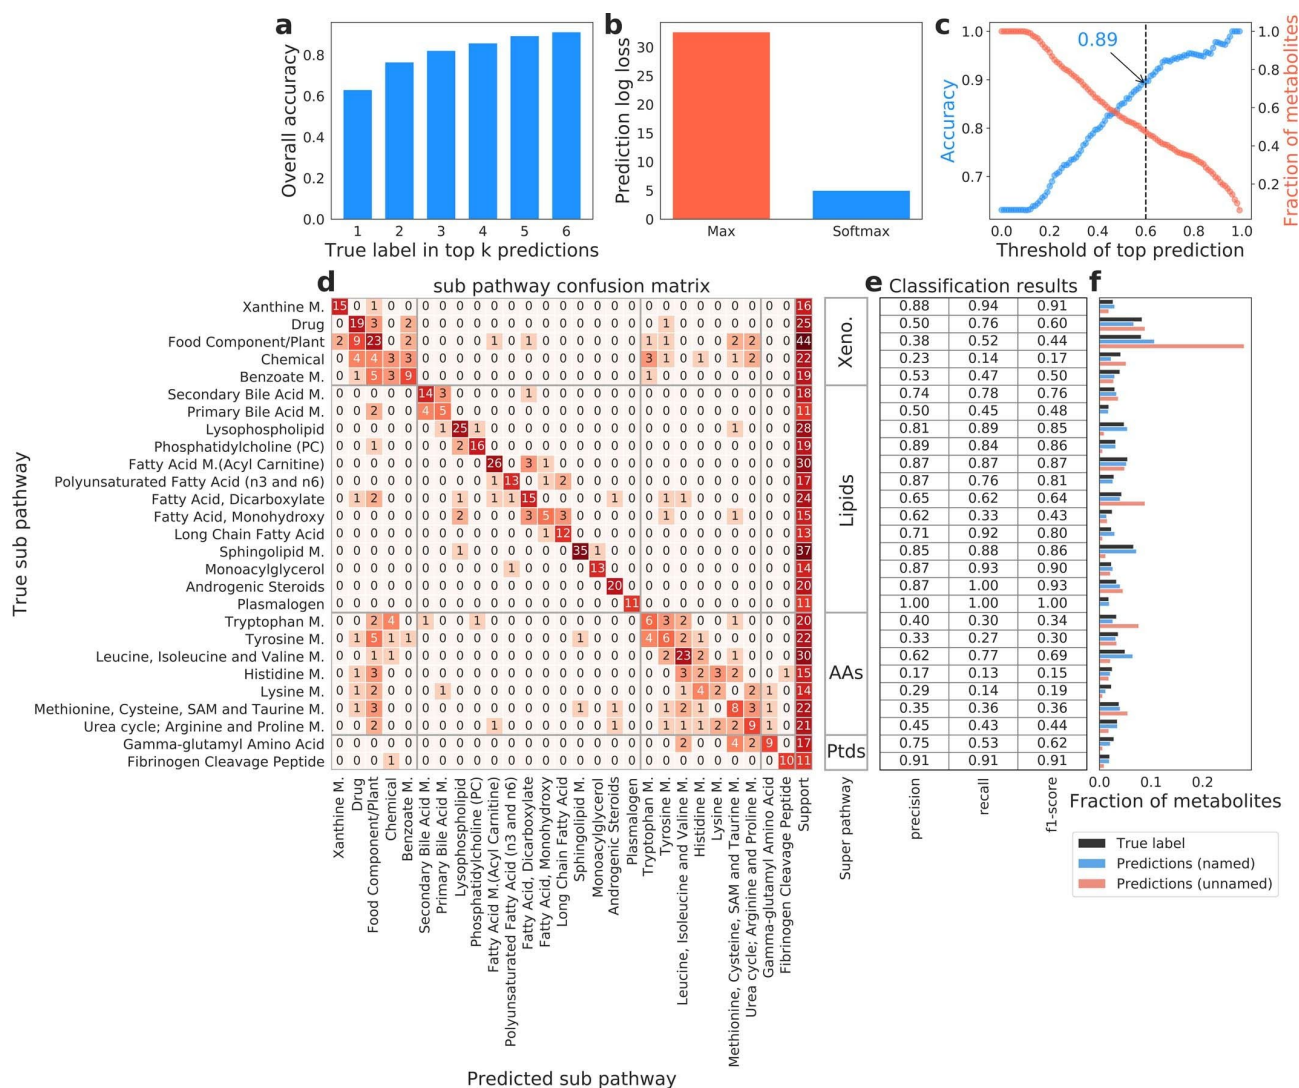
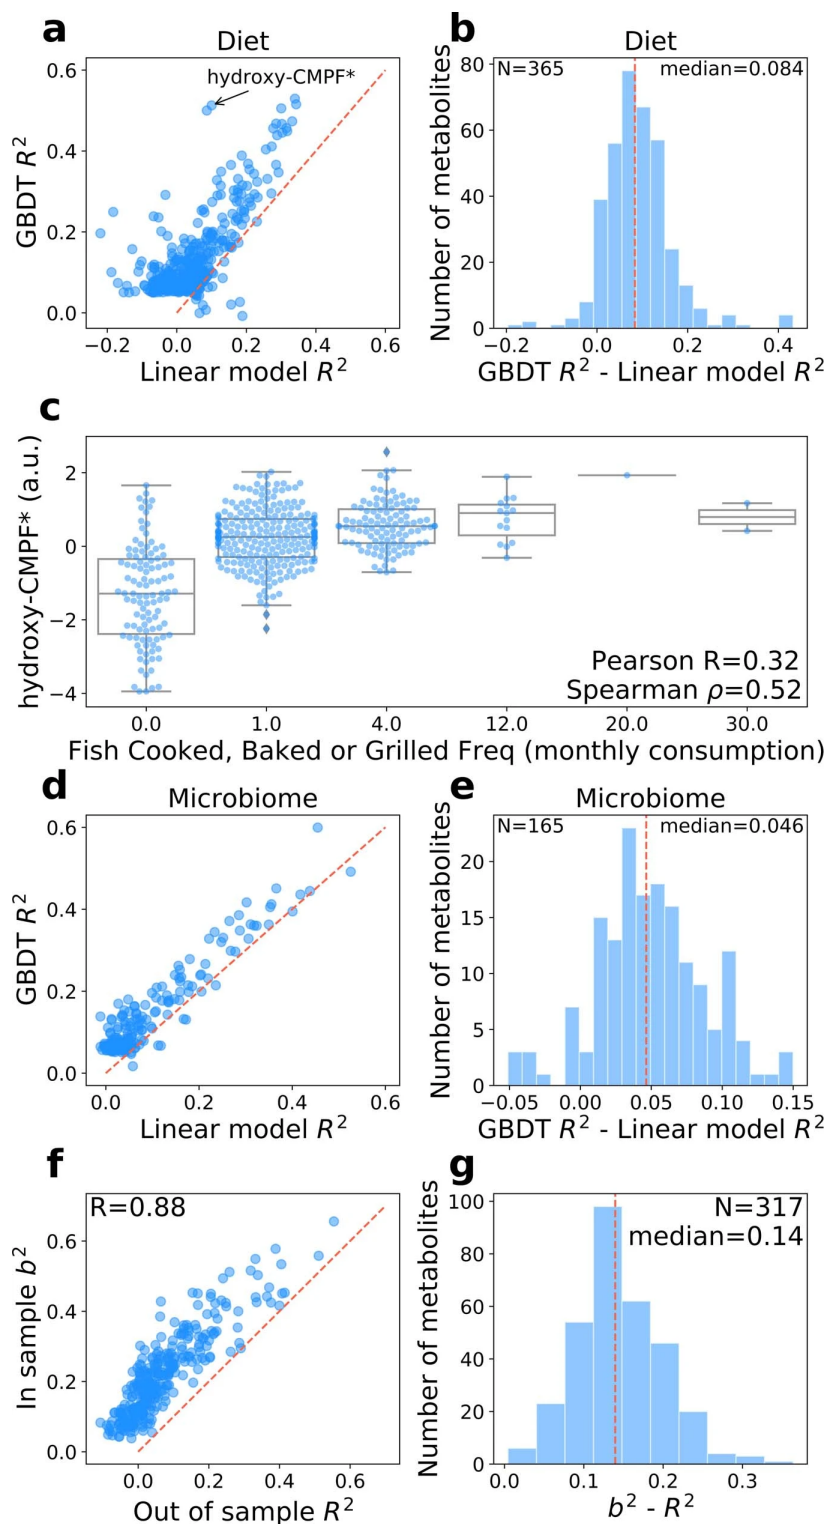


Figure panels refer to the results of a leave-one-out cross validation prediction model of sub pathways of metabolites based on their normalized levels, raw mean, percentage of missing values, and SHAP values (Methods). Results shown are for a model trained using only sub pathways that include over 10 molecules in our data (28 sub pathways, 572 named metabolites). **a**, The overall accuracy of the sub pathway classifier (y axis) when a success is considered as having the true label in one of the top k predictions (x axis). **b**, The log loss of the classifier (y axis) computed over the resulting soft max (raw probabilities; blue) and a dichotomous matrix in which for every metabolite we only keep the top predicted sub pathway as 1 and zero-out all other predictions (red). **c**, The overall accuracy of the model (left y axis; blue) and the corresponding fraction of metabolites (right y axis; red) when considering only metabolites for which the classifier predicted a maximal probability above some threshold (x axis). **d**, A confusion matrix showing the predicted sub pathways (x axis), determined as the label with the highest probability per metabolite, versus the true annotated sub pathways (y axis). Each cell in the matrix counts the number of metabolites of a certain true sub pathway (y axis) which were assigned with some predicted sub pathway (x axis) by our model. The rightmost column is the sum of every row and represents the number of metabolites annotated for every sub pathway. The matrix is ordered by the higher order biological pathway (super pathway). Cell colours are log scaled. **e**, Classification results summarizing the f1-score, precision and recall per sub pathway. Rows correspond to the sub pathway annotation in d. **f**, For

every sub pathway (y axis) shown are the fraction of metabolites truly annotated as such (black), predicted as such by the classifier (blue; out of the named molecules in the support of the model), and the fraction of unidentified molecules predicted as such (out of all unidentified molecules). M., metabolism; Xeno., xenobiotics; Ptds, peptides; AAs, amino acids.


Extended Data Fig. 3: Comparative analysis of linear versus nonlinear models and in-sample versus out-of-sample predictions.



a, Metabolite prediction R^2 of GBDT versus Lasso regression models using diet data. Shown are only metabolites for which at least one model achieved significant predictions with R^2 greater than 0.05. **b**, Histogram of the differences between the R^2 of GBDT compared to Lasso regression using the diet data. **c**, The levels of the metabolite hydroxy-CMPF* (y axis: centre, median; box, IQR; whiskers, $1.5 \times$ IQR) versus the monthly consumption of cooked, baked or grilled fish as reported in a food frequency questionnaire. The comparison of Spearman's and Pearson's correlation

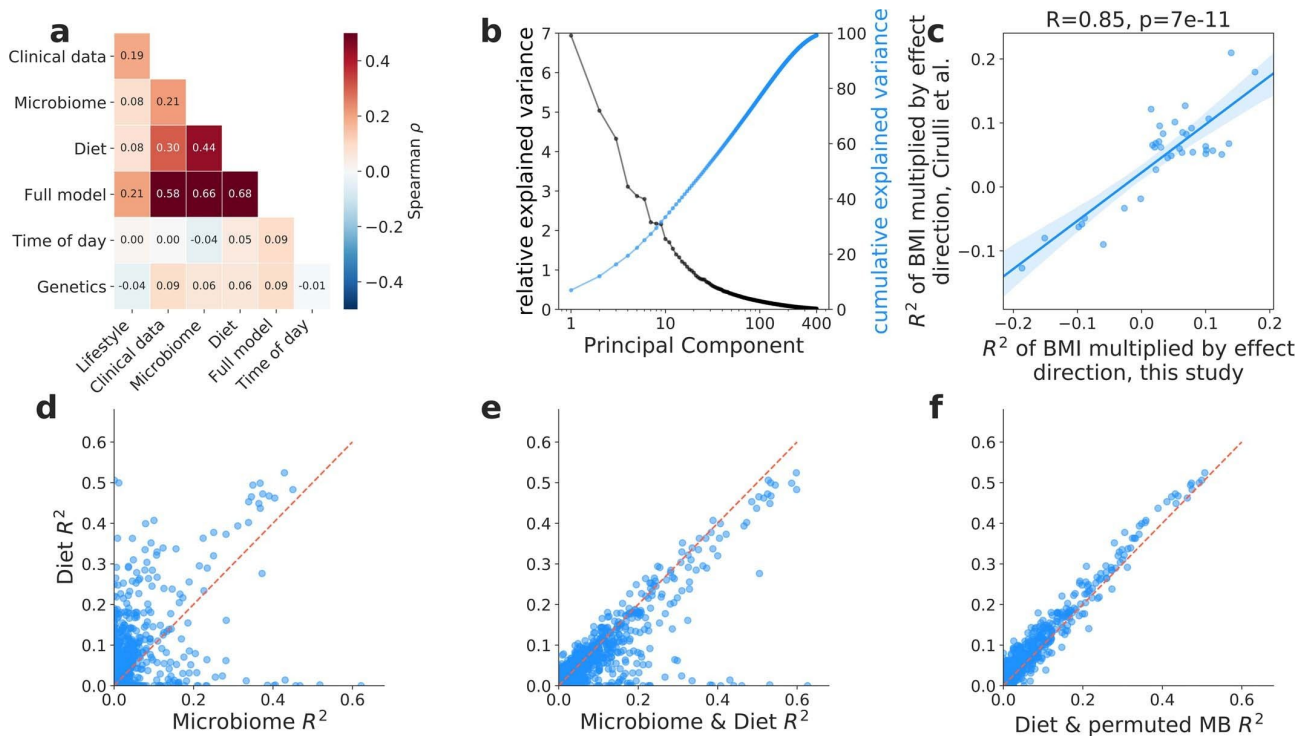
coefficients suggests that the relationship between the metabolite and the numerical values of the question are monotonic yet nonlinear, which explains why GBDT performs better in predicting the levels of hydroxy-CMPF* from diet data. The x axis is not to scale. **d, e**, Same as a, b for microbiome. **f**, Estimations of gut microbiome explainability (b_2) of metabolite levels obtained via applying a linear mixed model on the bacterial species composition as previously described (y axis) versus the explained variance (R^2) of metabolites from out-of-sample prediction models based on the same gut microbiome data. Shown are only metabolites with significant b_2 estimates (5% FDR). **g**, Histogram of the differences between the b_2 estimates and the R^2 of out-of-sample prediction using the gut microbiome data.

Extended Data Fig. 4: Comparison of explained variance of metabolites for every pair of feature groups.



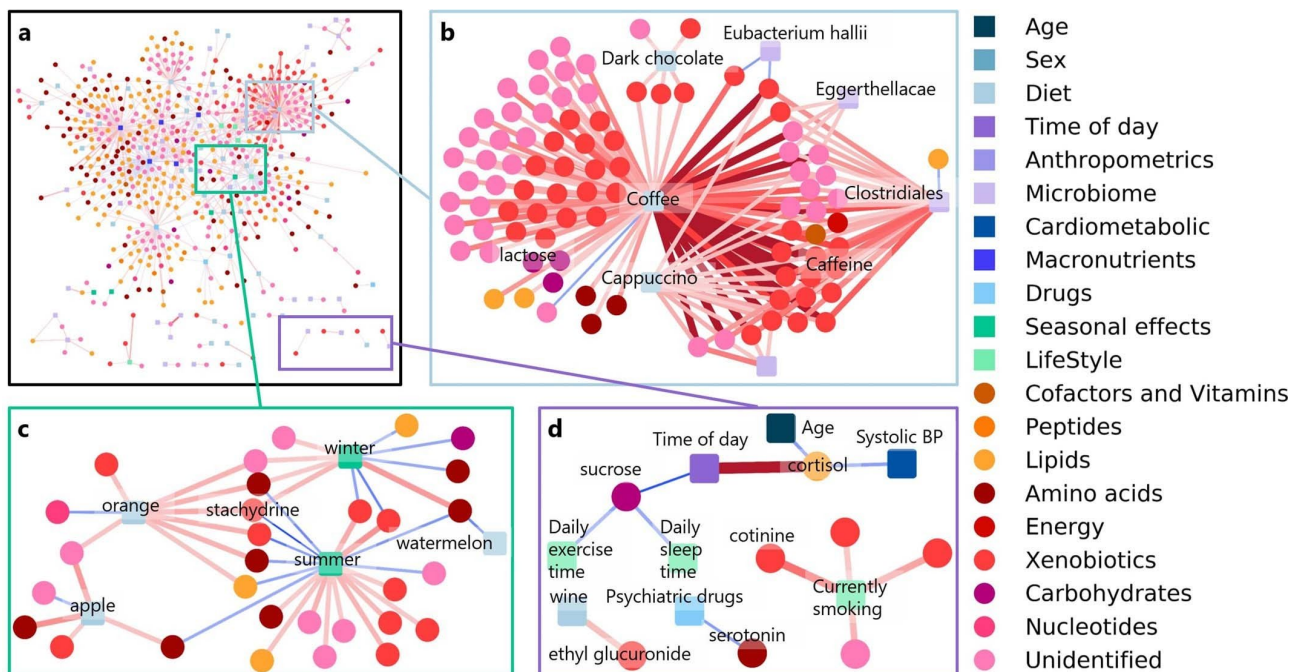
Dot plots of the explained variance of the metabolite groups (y axis) from models based on every pair of feature groups (x axis). Panels on the diagonal show the marginal distribution of explained variance of metabolite groups for a certain feature group.

Extended Data Fig. 5: Comparative analysis of different feature groups.



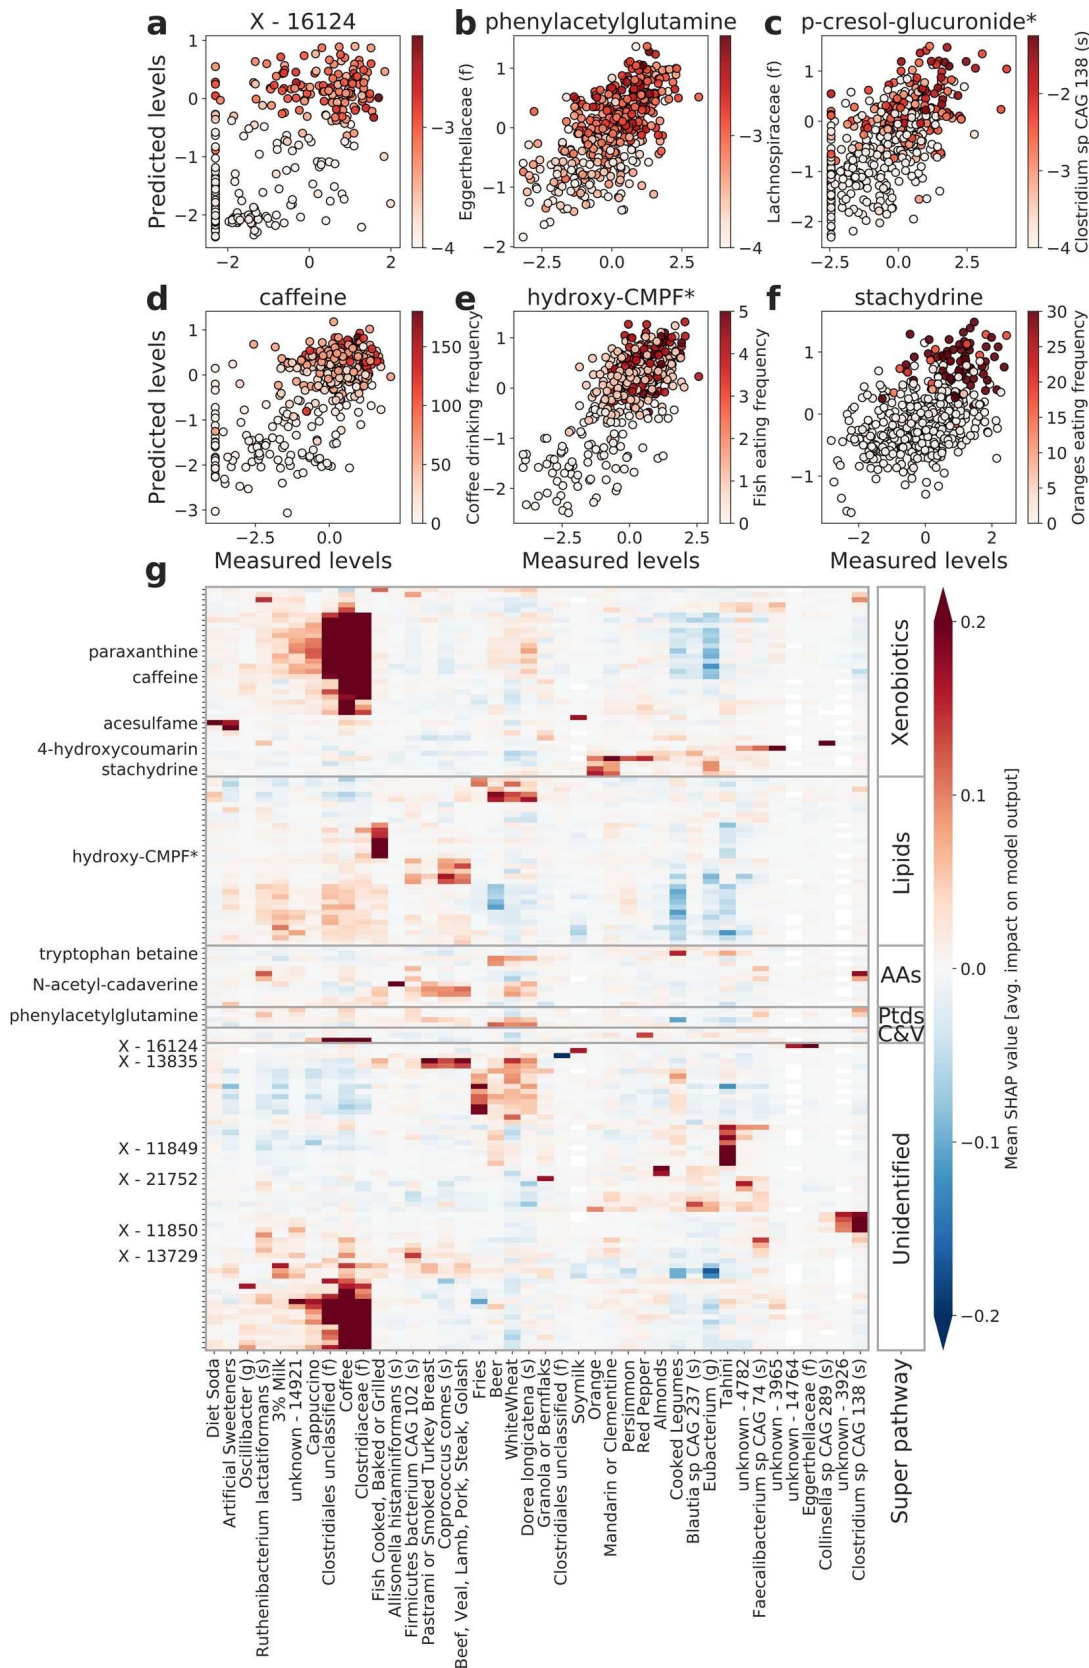
a, Spearman correlations computed between the EV of metabolites for every pair of feature groups. **b**, The proportion of variance explained by each of the first 400 principal components (left y axis; black) and their cumulative EV (right y axis; blue). **c**, R^2 multiplied by the sign of the Pearson correlation coefficient (x axis) between metabolite levels and BMI in our study, versus the mean R^2 multiplied by the sign of the Pearson correlation coefficient (y axis) of BMI associated metabolites recently reported by a different group¹³. Shown are 36 (out of 49) BMI associated metabolites that were also measured in this cohort. P value for the Pearson correlation, $P = 7 \times 10^{-11}$. Line and shaded colouring represent the fitting of a linear model and the 95% confidence interval. **d**, The EV of every metabolite from prediction models based on the gut microbiome (x axis) versus diet (y axis). Dashed red line is $y = x$. **e**, Same for prediction models based on both gut microbiome and diet (x axis) compared to using only diet (y axis). **f**, Same for prediction models based on diet and permuted gut microbiome (x axis) compared to using only diet (y axis). MB, microbiome.

Extended Data Fig. 6: Networks of interactions between phenotypes explain diverse metabolites.



Interactions between features from different feature groups predictive of similar metabolites are presented in a graphical layout, in which nodes are either metabolites or features, and edges are the directional mean absolute SHAP values (Methods) computed from models trained only on features from the respective feature group. Circular nodes, metabolites; predictive feature nodes, squares; both coloured by relevant categories. Shown are only edges with a mean absolute SHAP value greater than 0.12. **a**, Network of associations for the following feature groups: macronutrients, diet, microbiome, lifestyle, drugs and seasonal effects. **b**, A large group of metabolites for which predictions are mainly driven by the reported consumption of coffee and the relative abundance of a bacteria from the Clostridiales order. **c**, Metabolites explained by seasonal fruit consumption. **d**, Selected examples of interactions between metabolites and features in predictive models.

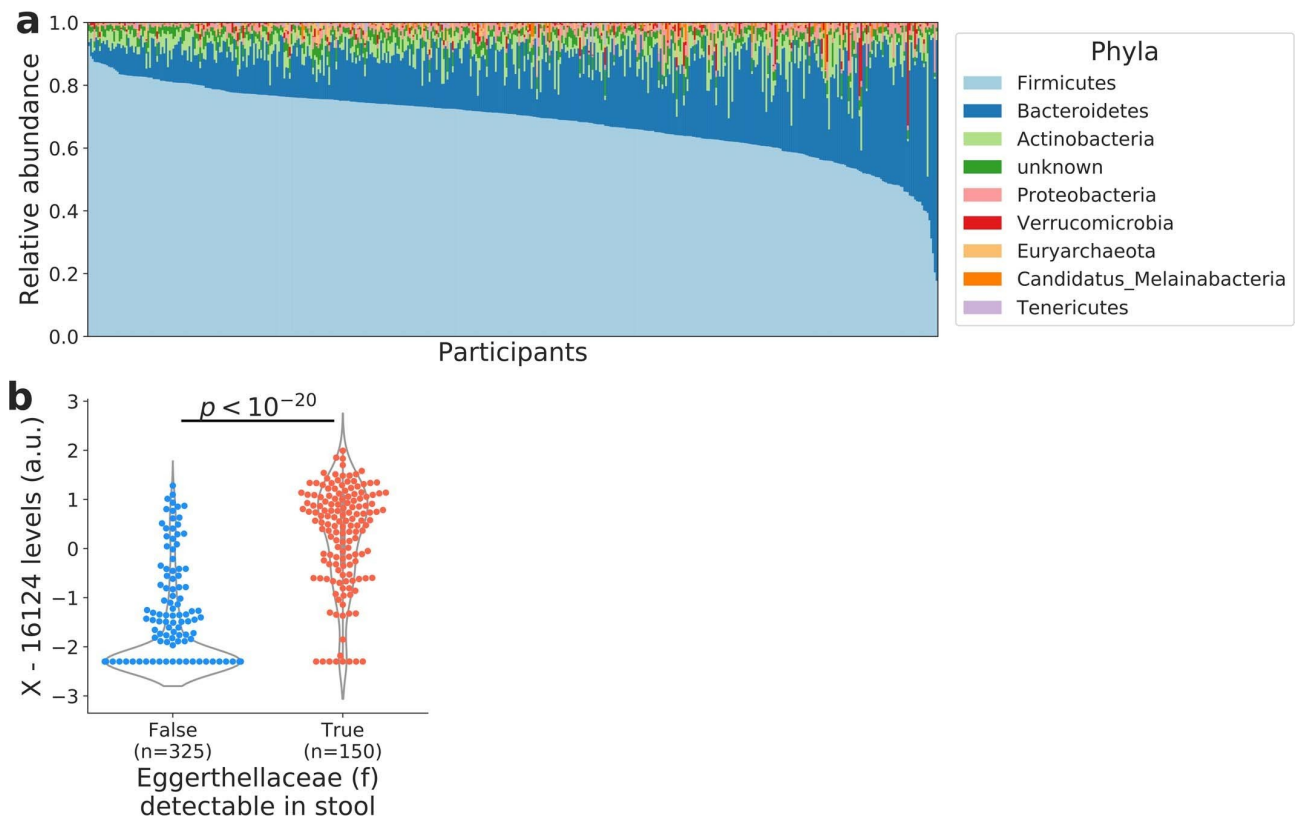
Extended Data Fig. 7: Specific dietary features and bacterial taxa underlie the accurate prediction of circulating metabolites.



a–f, Predicted (y axis) versus measured (x axis) levels (arbitrary units) of X-16124 (**a**; Pearson's $R = 0.77$, $P < 10^{-20}$), phenylacetylglutamine (**b**; $R = 0.63$, $P < 10^{-20}$), p-cresol-glucuronide (**c**; $R =$

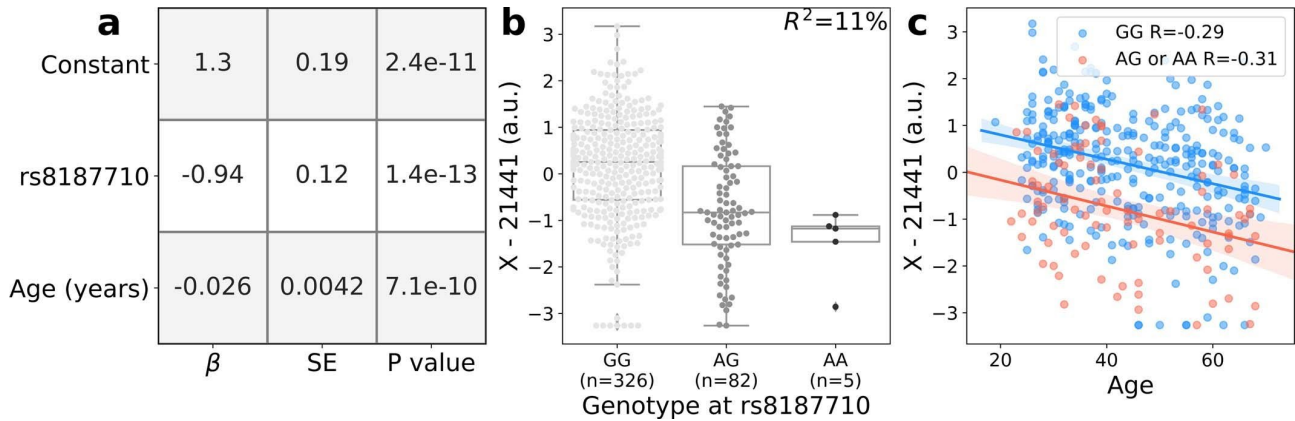
0.64, $P < 10^{-20}$), caffeine (**d**; $R = 0.68$, $P < 10^{-20}$), hydroxy-CMPF (**e**; $R = 0.72$, $P < 10^{-20}$) and stachydrine (**f**; $R = 0.5$, $P < 10^{-20}$). Predictions of a–c are based only on microbiome data, and coloured by the relative abundance of the bacterial taxa having the highest mean absolute SHAP value for each metabolite. Predictions of d–f are based only on diet data, and coloured by the reported consumption of the dietary item having the highest mean absolute SHAP value for each metabolite. P values for prediction were estimated via bootstrapping. **g**, Heat map showing the directional mean absolute SHAP values (Methods) of various features (x axis) computed from fivefold cross validation models that predict metabolite levels (y axis) using two separate models, one based on diet and another on gut microbiome data. Positive (negative) SHAP values indicate that higher (lower) feature values lead, on average, to higher predicted values. Shown are the top 150 predicted metabolites using diet and gut microbiome, and the top 40 features by maximum mean absolute SHAP value across all metabolites.

Extended Data Fig. 8: Distribution of phyla and a taxa from the Eggerthellaceae family.



a, Stacked bar plots per sample (x axis) showing the relative abundance of bacterial phyla (y axis). Samples are sorted by the relative abundance of the most abundant phylum, Firmicutes. Bacteroidetes is the second most abundant phylum in our cohort. Relative abundance of a phylum is computed as the sum over relative abundances of all bacterial features belonging to that phylum. **b**, The levels of the unidentified compound X-16124 in individuals for which the bacterial taxa from the Eggerthellaceae family was detectable in stool versus individuals for which it was not ($P < 10^{-20}$, two-sided Mann–Whitney U-test).

Extended Data Fig. 9: The unidentified molecule X-21441 associates with rs8187710 independent of age.



a, A table showing the coefficients, standard errors and P values resulted from a multiple linear regression model with levels of the unidentified molecule X-21441 as the dependent variable, the allele dosage of rs8187710 (0–2) and age (years) as the independent variables: $y_{X-21441} = \text{constant} + \beta_1 \times \text{rs8187710} + \beta_2 \times \text{Age}$. **b**, The levels of X-21441 (y axis; centre, median; box, IQR; whiskers, $1.5 \times \text{IQR}$) versus the genotype of the participants (x axis). The number of participants with each genotype is indicated below the tick labels. The explained variance of X-21441 by rs8187710 as estimated using plink (Methods) is indicated on the upper right corner of the panel. **c**, The levels of X-21441 versus the age of the participants (x axis) coloured by genotype of participants. Line and shaded colouring represent the fitting of a linear model and the 95% confidence interval. SE, standard error.

Extended Data Table 1 Basic characteristics and demographics of our main and validation cohorts

Characteristics	Main cohort (n=475)	TwinsUK (n=1004)	IMI-DIRECT (n=245)
Age (years)	44.0 +- 13.13	64.96 +- 7.77	61.81 +- 8.09
Sex, Males (%)	38.70%	3.90%	61.22%
BMI	26.3 +- 5.0	26.2 +- 4.71	30.59 +- 5.39
Smokers (%)	9.80%	42.40%	13.55%

Supplementary Information

Supplementary Note 1: Prediction of candidate biological pathways of unidentified compounds

In our data, unidentified molecules represent approximately 30% of the pool of currently identifiable metabolites in human serum. These unidentified molecules, being relatively unexplored, hold new opportunities for biomarker discovery and for a more complete understanding of the metabolic landscape. Therefore, a major priority is to infer their chemical identity and biological relevance. To that end, we developed a machine learning classifier to predict the biological pathway of metabolites based on their measured levels in our cohort, as well as on their interactions with the comprehensive profiling of our participants as captured by the SHAP values described above (Methods).

To evaluate our model, we used a leave-one-out cross validation scheme, in which we obtained a prediction of the high resolution biological pathway (termed sub pathway) for every named molecule. We found that the overall accuracy across subpathways which include over 10 molecules in our data was 63.1%. These pathways consist of the 28 most common subpathways, covering a total of 572 metabolites, where the fraction of the most common sub pathway was only 7.7%. A common evaluation metric for multi-label classification problems⁵³ is the accuracy of the model allowing for the true label to be among the top several predictions (usually 2-5). Having the true label among the few top predictions of a classifier implies that these next best predictions should also be considered when the uncertainty is high. Indeed, when allowing for the true label to be among the top two/five predictions, the overall accuracy went up to 76.6%/89.3% (**Extended Data Fig. 2a**), respectively, while the cumulative fraction of the two/five most common sub pathways was only 14.7%/30%. Accordingly, the log loss of the model's error was significantly lower when computed over the raw classifier output (probabilities), compared to the log loss computed over a dichotomous prediction vector (assigning 1 for the top scoring prediction, and 0 to all others; **Extended Data Fig. 2b**). For some sub pathways, mostly including structural related compounds such as plasmalogens, lysophospholipids but also for distinctive phenotype related compounds such as androgenic steroids and xanthine metabolites, our model achieved near perfect classification results, assigning almost all related metabolites with the right label. Yet, for several sub pathways, mostly classified as amino acids such as histidine and lysine metabolism, but also xenobiotics, for which the classification into drugs, chemicals or food/plant components is rather artificial and sometimes context dependent, our model had lower rates of precision and recall (**Extended Data Fig. 2d,e**). As another measure to estimate the performance and uncertainty of our model, we asked how its accuracy would change if we only considered cases in which the maximal predicted probability was above some threshold. We found that in cases for which our model assigned its top prediction with a probability of 0.6 or higher (covering 47.3% of the relevant metabolites), its overall accuracy was 89.7% (**Extended Data Fig. 2c**).

We next applied our model to the unidentified metabolites and obtained predictions of their underlying sub pathways (**Supplementary Table 2**). According to our model, the

unidentified molecules are significantly enriched for xenobiotics, as the fraction of compounds which the model predicted to be food or plant components (a sub pathway of xenobiotics) was 28%, compared to only 10.6% in the identified metabolites ($p < 10^{-10}$, Binomial test; **Extended Data Fig. 2f; Supplementary Table 3**). Other sub pathways which seem to be in higher abundance in the unidentified pool of molecules are tryptophan metabolism (7.6% vs 2.6%, $p < 10^{-5}$) and dicarboxylates (8.8% vs 4%, $p < 10^{-4}$). In contrast, our model predicts that some lipids such as sphingolipids and lysophospholipids are significantly underrepresented among the unidentified metabolites, which is further supported by the high recall of our model for these sub pathways. Still, we expect that the accuracy of the model for the unidentified molecules to be lower than for the identified ones, as the unidentified metabolites are likely to also cover less common sub pathways which are not represented in the pool of named molecules. As partial validation of the performance of our model on the unidentified molecules, we examined its biological pathway predictions on a set of 85 unidentified metabolites which were recently identified or partly identified by Metabolon (**Supplementary Table 4; Methods**). Among these 85 metabolites, only 43 belong to sub pathways which were covered by our model, while the rest are either partly characterized (17), belong to an uncommon pathway (8) or were only assigned with a candidate structure (17). Using only our model's best guess, we correctly predicted the sub pathway for 18 out of the 43 metabolites, while the true label was among the top three predictions in 30 out of 43 cases (**Supplementary Table 5**).

Of particular interest are metabolites which are accurately predicted by the gut microbiome, as they may be modulated by perturbing the bacterial community. Since many of the metabolites that were predicted by the gut microbiome with high accuracy are unidentified, our predictions of their biological pathway may help elucidate their role in chemical reactions and in bacterial metabolism in general. Among these molecules, some are predicted by the microbiome with the highest accuracy, including X-11850, X-12126, X-12216 and X-11843. These were all predicted with $R^2 > 0.25$ using the microbiome, and according to a recent effort by Metabolon, are likely to be derivatives of aromatic amino acids, a class of molecules known to be metabolized by the gut microbiome⁵⁴. Indeed, our model predicts that these metabolites are most likely to take part in the tryptophan and tyrosine metabolism pathways. Another prominent class of molecules are secondary bile acids derivatives, which are derived by the enzymatic action of intestinal bacteria^{55,56}. Notable examples of these molecules include X-17469, X-14662, X-16654 and X-11491, which Metabolon has recently identified as lithocholic acid sulfate, glycooursodeoxycholate sulfate, deoxycholic acid 12-sulfate and deoxycholic acid glucuronide, with accordance to our model's prediction of these metabolites being secondary bile acid derivatives. Using our model, we predict that several other unidentified metabolites are in fact part of the secondary bile acid metabolism pathway, such as X-22834, X-22520, X-17612, X-19438 and X-21467.

Overall, our model persistently predicts that among the 75 unidentified metabolites best explained by gut microbiome data ($R^2 > 5\%$), most are xenobiotics (e.g. food or plant components, benzoate metabolites and other chemicals), aromatic amino acid derivatives (tryptophan, histidine and tyrosine metabolites), fatty acids and secondary bile acid derivatives. Finally, as unidentified molecules constitute a large portion of metabolites explained by the human microbiome, our predictions of their biological pathways constitute a major step towards mapping the metabolic producing and modulating potential of the human gut microbiome.

Supplementary Note 2: Out-of-sample performance evaluation

Across this study we purposely decided to report only estimates computed via out-of-sample predictions, since we view such evaluation of performance based only on unseen samples as the most strict and conservative estimate of performance. As such, our results constitute a very conservative estimate for the amount of variance in metabolite levels that may be explained by the various features we examined. Furthermore, unlike age, sex, or other clinical parameters which have a simple structure and can be measured accurately, features such as diet and gut microbiome composition have inherent measurement noise and their complexity makes them much harder to represent, further emphasizing that our results are conservative, as advances in methodologies for analysis of such data may increase prediction power. The heterogeneity of our data is yet another aspect in which out-of-sample prediction comes in handy, since its estimates do not depend on modeling assumptions, as compared to in-sample methodologies such as linear mixed models (LMM) which are commonly used in estimation of genetic heritability⁵⁷. Indeed, by applying a LMM-based methodology that we have recently proposed for the estimation of explainability of the gut microbiome composition³³, we demonstrate that estimates obtained via out-of-sample predictions are much more conservative, as we can obtain a median gain of 13.9% in EV of metabolites levels by applying these in-sample models (**Extended Data Fig. 3f,g**; Methods).

Supplementary Note 3: Features interact in metabolome predictions

We next examined how different feature groups interact in explaining the levels of various metabolites, by building separate models based on each, and using SHAP in order to estimate the impact of each feature on the output of the models. We uncovered a dense network of interactions between feature groups in explaining metabolite levels (**Extended Data Fig. 6a**; Methods).

As mentioned above, we found that coffee consumption was linked to a large number of metabolites, most of which are unidentified compounds and xenobiotics from the xanthine metabolism pathway. Notably, we found that a specific Clostridiales species was linked to many of these (**Extended Data Fig. 6b**), suggesting a possible interaction between coffee consumption and the presence of these bacteria in explaining the levels of these metabolites. Being the most predictive features among their feature categories, coffee consumption and this Clostridiales species may be targets for validation using interventional studies.

The consumption of citrus fruits such as oranges positively affected (on average) the prediction of several metabolites such as stachydrine, a known biomarker for the consumption of citrus fruits⁵⁸ (also named proline betaine; significantly predicted by diet, Pearson $R=0.50$, $p<10^{-20}$), which in turn had higher values in samples taken in winter months, consistent with the fact that oranges are seasonal fruits available in Israel mostly during winter (**Extended Data Fig. 6c**). Another example is N-methyltaurine ($R=0.35$, $p<10^{-20}$), an amino acid which has higher levels in samples taken during winter, and whose prediction was negatively affected by the consumption of watermelon, a summer seasonal fruit (**Extended Data Fig. 6c**).

Additionally, to further assess the quality of our results, we explored several associations which were previously reported (**Extended Data Fig. 6d**). We found that the levels of cortisol were lower in samples taken during the second half of the day (Prediction with time of day, $R=0.63$,

$p < 10^{-20}$, positive SHAP value for samples taken in the morning), consistent with previous studies showing that cortisol levels peak early in the morning⁵⁹. We also found that tobacco-related metabolites such as cotinine (Prediction $R=0.72$ by lifestyle, $p < 10^{-20}$) had higher levels in active smokers (positive SHAP values for smoking), and that no other feature could significantly explain their levels. Finally, we found that blood serotonin levels (Prediction $R=0.46$ by drugs, $p < 10^{-6}$) were lower in participants who reported taking psychiatric drugs (negative SHAP values), despite serotonin being a therapeutic target for selective serotonin reuptake inhibitors⁶⁰ prescribed to increase serotonin levels in the brain.

Supplementary Note 4: Limited statistical power for the identification of low effect size genetic variants

We previously argued³³ that assessing the impact of human genetics on continuous traits (SNP heritability) requires much larger cohorts, as high confidence heritability estimates are feasible only when cohort size exceeds several thousands of participants. We nevertheless included analysis of human genetics in this study, as it is well established that genetics have a wide influence over the levels of blood metabolites^{2,3,21-24}. As the size of our cohort is moderately small compared to other cohorts used in the field of human genetics, ranging between tens of thousands to hundreds of thousands⁶¹, we were limited in terms of the statistical power needed for the identification of relatively small effect variants ($R^2 < 10\%$), hence we only report of associations with large effect size. Reporting only large effect variants further introduces bias to the estimates, as we were unable to include associations of smaller effect size, which are highly prevalent and likely to occur, even if they are detectable in our data. Though our cohort is not the largest in which serum metabolomics was coupled with human genetics, it is still the only cohort in which serum metabolomics was coupled with other phenotyping assays, allowing comparison of the relative contribution of many established potential determinants of the serum metabolome. It will be interesting to revisit these comparisons of the predictive power of genetics, diet, and gut microbiome and validate our results in a much larger cohort in which all these data are profiled.

References

53. Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. in *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255 (IEEE, 2009). doi:10.1109/CVPR.2009.5206848
54. Dodd, D. *et al.* A gut bacterial pathway metabolizes aromatic amino acids into nine circulating metabolites. *Nature* **551**, 648–652 (2017).
55. Begley, M., Gahan, C. G. M. & Hill, C. The interaction between bacteria and bile. *FEMS Microbiol. Rev.* **29**, 625–651 (2005).
56. Ridlon, J. M., Kang, D. J., Hylemon, P. B. & Bajaj, J. S. Bile acids and the gut microbiome. *Curr. Opin. Gastroenterol.* **30**, 332–338 (2014).
57. Yang, J. *et al.* Common SNPs explain a large proportion of the heritability for human height. *Nat. Genet.* **42**, 565–569 (2010).
58. Atkinson, W., Downer, P., Lever, M., Chambers, S. T. & George, P. M. Effects of orange juice and proline betaine on glycine betaine and homocysteine in healthy male subjects. *Eur. J. Nutr.* **46**, 446–452 (2007).
59. Smyth, J. M. *et al.* Individual differences in the diurnal cycle of cortisol. *Psychoneuroendocrinology* **22**, 89–105 (1997).
60. Hyttel, J. Pharmacological characterization of selective serotonin reuptake inhibitors (SSRIs). *Int. Clin. Psychopharmacol.* **9 Suppl 1**, 19–26 (1994).
61. Sudlow, C. *et al.* UK biobank: an open access resource for identifying the causes of a wide range of complex diseases of middle and old age. *PLoS Med.* **12**, e1001779 (2015).