4 August 2025

Profiling Biological Effects of Microbiome Metabolites via Machine Learning

Hong A. Chung[1], Zachary Fralish[2], Tiffany Tu[2], Daniel Reker[2]

1. Duke

2. Duke University

Abstract

Human microbiome-derived metabolites are key mediators of host physiology. However, their biological effects remain largely uncharacterized due to limitations of current low-throughput and untargeted experimental approaches that are time-intensive and costly. This has hindered the systematic biological characterization of microbiome metabolites. To address this gap and accelerate the identification of biological effects of microbiome metabolites, we developed a machine learning platform trained on publicly available drug development data to rapidly predict a wide array of chemical and biological properties of microbiome metabolites. Prospective experimental validation confirmed the predictive accuracy of our models and uncovered previously unknown effects of several metabolites. For example, we identified previously unknown Interleukin-8 secretion stimulation by the metabolites spermine and spermidine, which have been regarded anti-inflammatory thus far. Our findings demonstrate the power of machine learning to accelerate the functional annotations of microbiome-derived metabolites, paving the way for new biomarker discovery and therapeutic development.

Keywords

Microbiome metabolites, machine learning, ADMET, Interleukin-8 secretion

Profiling Biological Effects of Microbiome Metabolites via Machine Learning

Hong A. Chung[1], Zachary Fralish[1], Tiffany Tu[2], Daniel Reker[1, 3, *]

[1] Department of Biomedical Engineering, Duke University, NC 27708, USA

[2] Department of Computational Biology and Bioinformatics, Duke University, NC 27708, USA

[3] Duke Microbiome Center, Duke University, Durham, NC 27710, USA

[*] Corresponding Author: Daniel Reker, daniel.reker@duke.edu

ABSTRACT

Human microbiome-derived metabolites are key mediators of host physiology. However, their biological effects remain largely uncharacterized due to limitations of current low-throughput and untargeted experimental approaches that are time-intensive and costly. This has hindered the systematic biological characterization of microbiome metabolites. To address this gap and accelerate the identification of biological effects of microbiome metabolites, we developed a machine learning platform trained on publicly available drug development data to rapidly predict a wide array of chemical and biological properties of microbiome metabolites. Prospective experimental validation confirmed the predictive accuracy of our models and uncovered previously unknown effects of several metabolites. For example, we identified previously unknown Interleukin-8 secretion stimulation by the metabolites spermine and spermidine, which have been regarded anti-inflammatory thus far. Our findings demonstrate the power of machine learning to accelerate the functional annotations of microbiome-derived metabolites, paving the way for new biomarker discovery and therapeutic development.

KEYWORDS

Microbiome metabolites, machine learning, ADMET, Interleukin-8 secretion

INTRODUCTION

The human microbiome produces thousands of metabolites that can influence host health and disease[1] [2] [3]. For example, microbiome-derived tryptophan has been associated with modulation of the liver's inflammatory responses[4] and microbiome-derived inosine has been found to exhibit immunomodulatory effects that can support immunotherapeutic responses in mice[5]. Conversely, microbiome-derived trimethylamine n-oxide is under scrutiny for its potential role in cardiovascular disease[6]. Such microbiome-host interactions are considered highly medically relevant since these associations could ultimately lead to the discovery of important biomarkers and serve as informed starting points for the development of novel therapeutic interventions. However, only a minute fraction of microbiome metabolites has been characterized for their biological effects[7]. This is primarily driven by the cost and complexity of top-down analysis of microbiome health effects, which largely rely on laborious correlative studies using metagenomic or metabolomic data with comparably low throughput, limited experimental resolution, and often small sample sizes[8].

Recent advances in machine learning offer a path to bypass these constraints as substantial progress has been made in using machine learning for natural product-based drug discovery[9] [10]. Bespoke computational tools can now anticipate the potential biological effects of metabolites from various organisms including microbes and plants. This is achieved using increasingly accurate algorithms that process large datasets to associate the chemical structures of metabolites with those of molecules with known biological effects, enabling the prediction of similar bioactivity. These tools allow drug hunters in industry and academia to triage their experimental testing, thereby saving valuable resources compared to alternative approaches such as high-throughput screening of natural product libraries or affinity-guided purification approaches[11] .

Here, we propose that similar computational approaches could accelerate the identification of biological effects of microbiome metabolites (**Fig. 1**). By building on publicly available data repositories of millions of small molecules with known biological effects and properties, we train machine learning models to predict a wide range of biological properties for microbiome metabolites including permeability across various biological barriers, liver toxicity, or immune stimulation through activation of Interleukin-8 (IL-8) secretion.

In this study, we define "microbiome metabolites" as compounds that can be biosynthesized or chemically modified by microbial metabolism. We base our knowledge on the origin and the modification of metabolites on source annotations from established, peer-reviewed public databases[12] [13] [14] [15] [16]. We acknowledge that source annotations vary in quality and completeness across databases and may not fully capture biosynthetic origins. As such, this classification serves as an operational framework rather than a strict biological definition. To ensure rigor, any top candidate identified through model predictions and considered for downstream applications should undergo manual review of biosynthetic context to confirm its microbial origin (**Table S1**). Additionally, we do not exclude metabolites that are also produced by other organisms, including the human host, since its production by microbiome might nevertheless alter exposure and local concentrations of metabolites in physiological conditions[17].

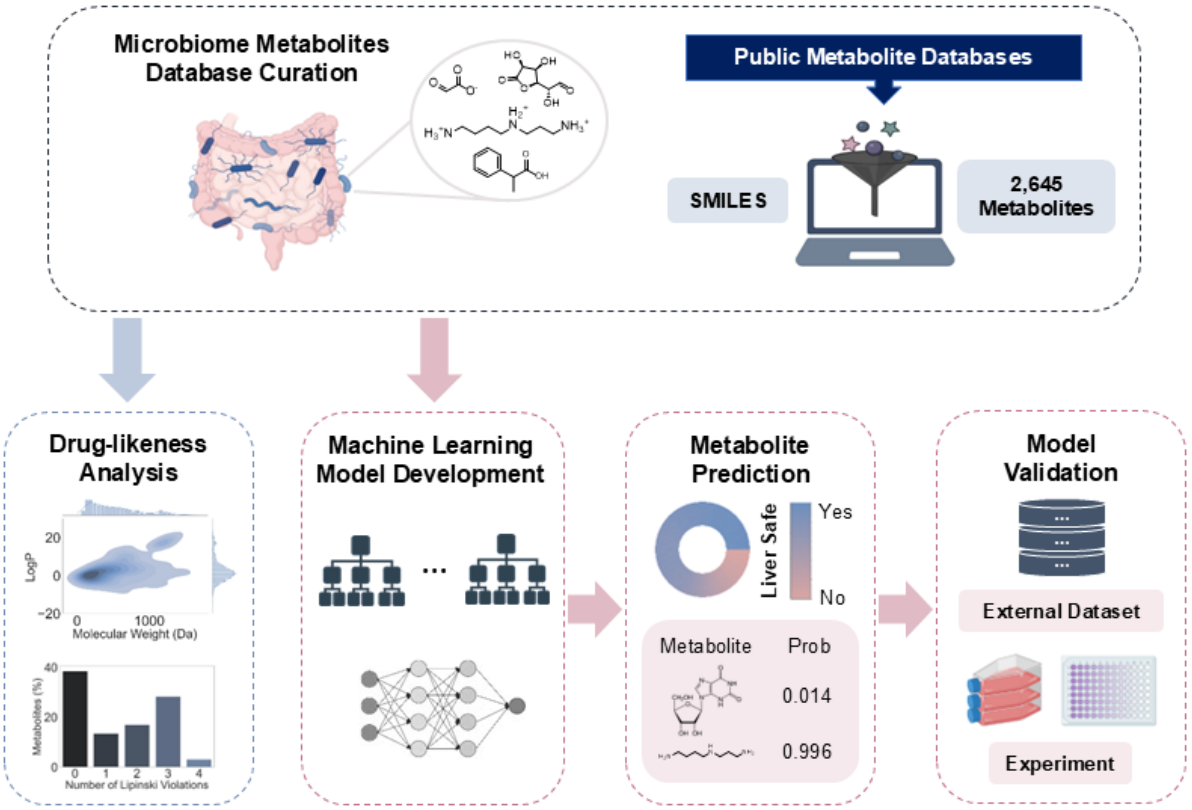

Figure 1. General workflow of the presented study. A database of human microbiome metabolites was semi-automatically curated by mining data from five different public databases to extract chemical structures (SMILES) and the microbial origin of the metabolites. The physicochemical properties of these metabolites are calculated using cheminformatics tools (RDKit). More complex biological properties were predicted using machine learning models (KNIME data analytics platform) trained on data for drug-like molecules (TDCommons or PubChem). We used external datasets and conducted prospective *in vitro* experiments to validate our models.

To assess the utility of our models, we evaluated the performance of our models on external test data and conducted prospective *in vitro* experiments in-house. While performance varied across different tasks, we consistently found strong performance of our models with an average accuracy of 81.7% on external validation sets and an average accuracy of 92.86% in prospective experiments. In addition, our experiments identified previously unknown properties of specific microbiome metabolites, such as the IL-8 stimulatory potential of the metabolites spermine and spermidine.

Taken together, we here validated that machine learning can be deployed to triage the identification of biological effects of microbiome metabolites. These models can prioritize candidate metabolites for *in vitro* experimental testing, thereby shifting microbiome metabolite research from explorative and expensive screening experiments to targeted validation studies. We expect that this approach and similar platforms could accelerate microbiome research by efficiently identifying metabolites with specific biological effects that could in the future serve as biomarkers or as starting points for the development of novel therapeutic interventions.

RESULTS & DISCUSSION

Microbiome metabolites have "drug-like" properties

To ensure that machine learning models trained on drug-like molecules could be applicable to microbiome metabolites, the metabolites need to share chemical features with the drug-like molecules in the training data to allow algorithms to recognize learned molecular patterns and make accurate predictions for metabolites[18]. As a first simple assessment of whether microbiome metabolites share similarities with drug-like molecules, we generated a t-distributed stochastic neighbor embedding (t-SNE) as a dimensionality reduction to visually assess the neighborhood relationships in chemical structure spaces occupied by 2,645 microbiome metabolites and 1,855 FDA-approved drugs (**Fig. 2A**). We observed a notable overlap between the metabolites and approved drugs in chemical space, suggesting that several metabolites share molecular patterns with drug-like molecules. This implies that models trained on drug-like structures might be applicable to microbiome metabolites. This finding also hints at many metabolites having chemical structures that could enable them to modulate biological pathways in hosts similarly to approved medicines, as they were identified as "drug-like". Notwithstanding this notable overlap, we did observe that there were also regions in chemical space that were exclusively occupied by metabolites or exclusively occupied by drugs, suggesting that certain types of drugs or metabolites are chemically distinct from the other group. For example, we observed a large group of metabolites that formed a distinct cluster that did not resemble any drug-like structures (**Fig. 2A**), which upon further inspection was largely composed of lipids - key metabolites involved in various biological functions, including membrane formation, energy storage, and cell signaling, but chemically distinct from currently approved drug-like molecules.

Beyond chemical structure similarity, we also aimed to assess whether microbiome metabolites and drug-like structures exert similar basic physicochemical properties. We analyzed the distribution of six different molecular properties (**Fig. 2B**), including the logarithm of the molecule's partition coefficient (logP), the molecular weight, the number of hydrogen-bond donors, the number of hydrogen-bond acceptors, the number of rotatable bonds, and the molecular topological polar surface area (TPSA). We observed similar property distributions between drugs and microbiome metabolites for all these properties (**Fig. 2B**). Additionally, for further contextualization, we calculated these molecular properties for metabolites from the *Pseudomonas aeruginosa* Metabolome Database (PAMDB) [19]. *P. aeruginosa* differs from typical commensal microbes as it is a pathogenic bacterium known to cause infections in the blood, lungs, or urinary tract in humans [20]. We observed that the *P. aeruginosa* metabolites differed in their property distributions compared to drugs and microbiome metabolites (**Fig. 2B**). Although we cannot discount the influence of reporting biases on these observed differences, this analysis continues to highlight a considerable similarity between many microbiome metabolites and drug-like compounds that can potentially be exploited by machine learning algorithms trained on drug-like structures to predict metabolite activities.

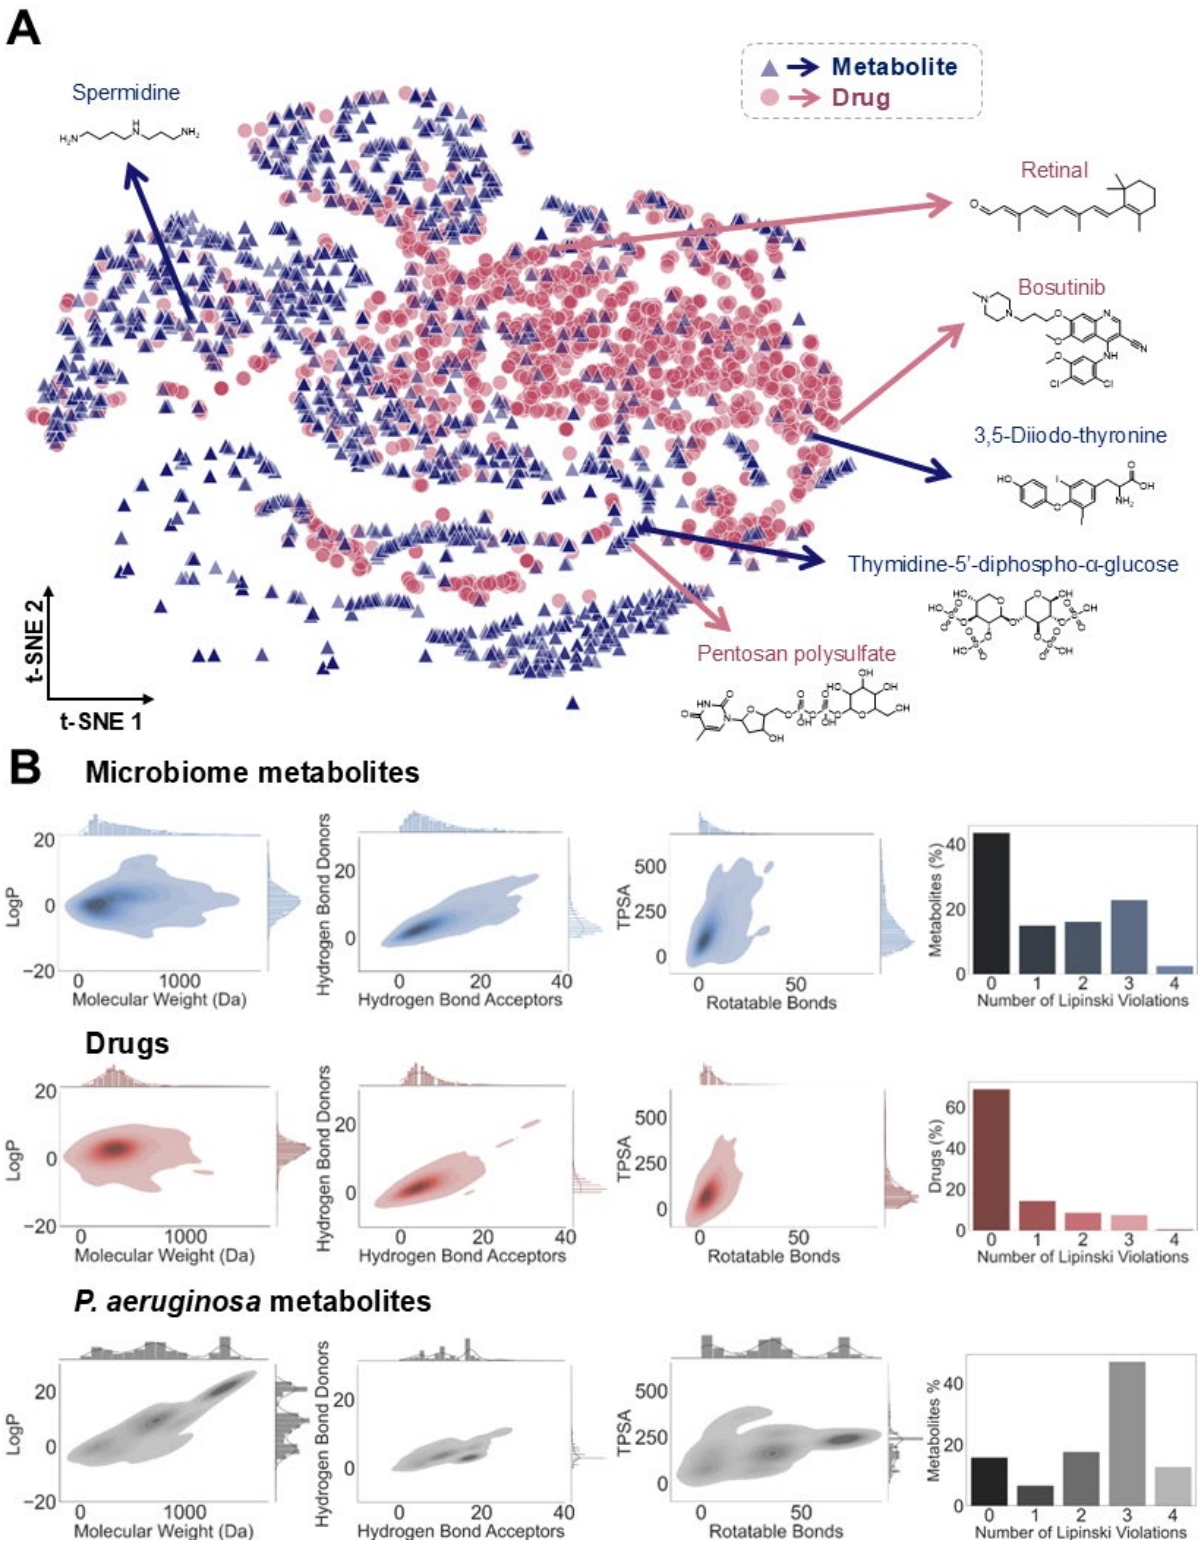

Figure 2. Microbiome metabolites are drug-like. (**A**) Chemical space of microbiome metabolites and FDA-approved drugs. A t-SNE plot was generated using Morgan fingerprints and RDKit descriptors in KNIME. This dimensionality reduction reveals that many metabolites are "drug-like", as they overlap with approved drugs in chemical space. Chemical structures of a few representative microbial metabolites (blue triangles) and drugs (red circles) are shown. (**B**) Six chemical properties were calculated and analyzed for human microbiome metabolites, FDA-approved drugs, and metabolites of the pathogenic microbe *P. aeruginosa*: LogP (logarithm of the partition coefficient), molecular weight, number of hydrogen bond donors, number of hydrogen bond acceptors, topological polar surface area (TPSA), and the number of rotatable bonds. These properties were then evaluated against Lipinski's "Rule of Five" criteria. According to Lipinski's "Rule of Five", more than 40% of the microbial metabolites meet all the four criteria to be considered "drug-like".

The first four of the investigated chemical properties (logP, molecular weight, number of hydrogen bond doners, number of hydrogen bond acceptors) are also commonly used to determine whether molecules are compliant with Lipinski's "Rule of Five". Originally, Lipinski's "Rule of Five" had been developed to serve as a simple heuristic guideline for gauging the oral bioavailability of a synthetic small molecule compound[21]. While no longer considered an accurate filter due to the recognition of important exceptions, these rules remain utilized as indicative benchmarks for properties crucial to a molecule's bioavailability and their "drug-likeness" [22]. We found that 43.55% of the microbiome metabolites meet all four criteria stipulated by Lipinski's "Rules of Five" and 14.98% have only one violation. This indicates that more than half (58.54%) of the microbiome metabolites are compliant with Lipinski's "Rules of Five" by not violating more than one of the rules. In contrast, only 15.83% of the *P. aeruginosa* metabolites meet all four criteria and 6.65% have only one violation – indicating that metabolites produced by other microbes might be less "drug-like" compared to human microbiome metabolites.

Overall, this analysis suggests that microbiome metabolites share similar molecular patterns and properties with drug-like structures, which indicates that machine learning models trained on drug-like structures may be well suited for guided elucidation of properties and biological effects of microbiome metabolites.

Machine learning model development for ADMET property prediction

Since our chemical similarity assessment and physicochemical property comparisons indicated that many microbiome metabolites might be sufficiently similar to approved drugs and thereby could be processed by molecular machine learning algorithms[23], we next aimed to predict biological properties of the microbiome metabolites with machine learning. For this, we first focused on various absorption, distribution, metabolism, excretion, and toxicity (ADMET) properties of microbiome metabolites. We trained machine learning classification models on publicly available datasets of small molecules for various ADMET properties[24] (**Table S2**) to predict classes of microbiome metabolites in terms of these properties (**Fig. 3A**).

First, to identify the most promising predictive machine learning models for each of the properties of interest, we systematically tested five different models for each property in retrospective cross-validation experiments. We explored different predictive architectures ranging from classical machine learning models such as Random Forest to deep neural networks such as multilayer

perceptrons (MLP). For the ADMET properties of interest, the XGBoost Tree (XGBT) model consistently outperforms all other tested models, followed by either Random Forest (RF) or Support Vector Machine (SVM), when evaluating the model's predictive ability using the Cohen's κ metric (**Fig. S1**). Cohen's κ measures how well a classifier predicts the correct class of test set datapoints compared to agreement by random chance contextualizing classifier performance, especially in situations where the data is highly imbalanced, since random agreement can be more likely in imbalanced scenarios.

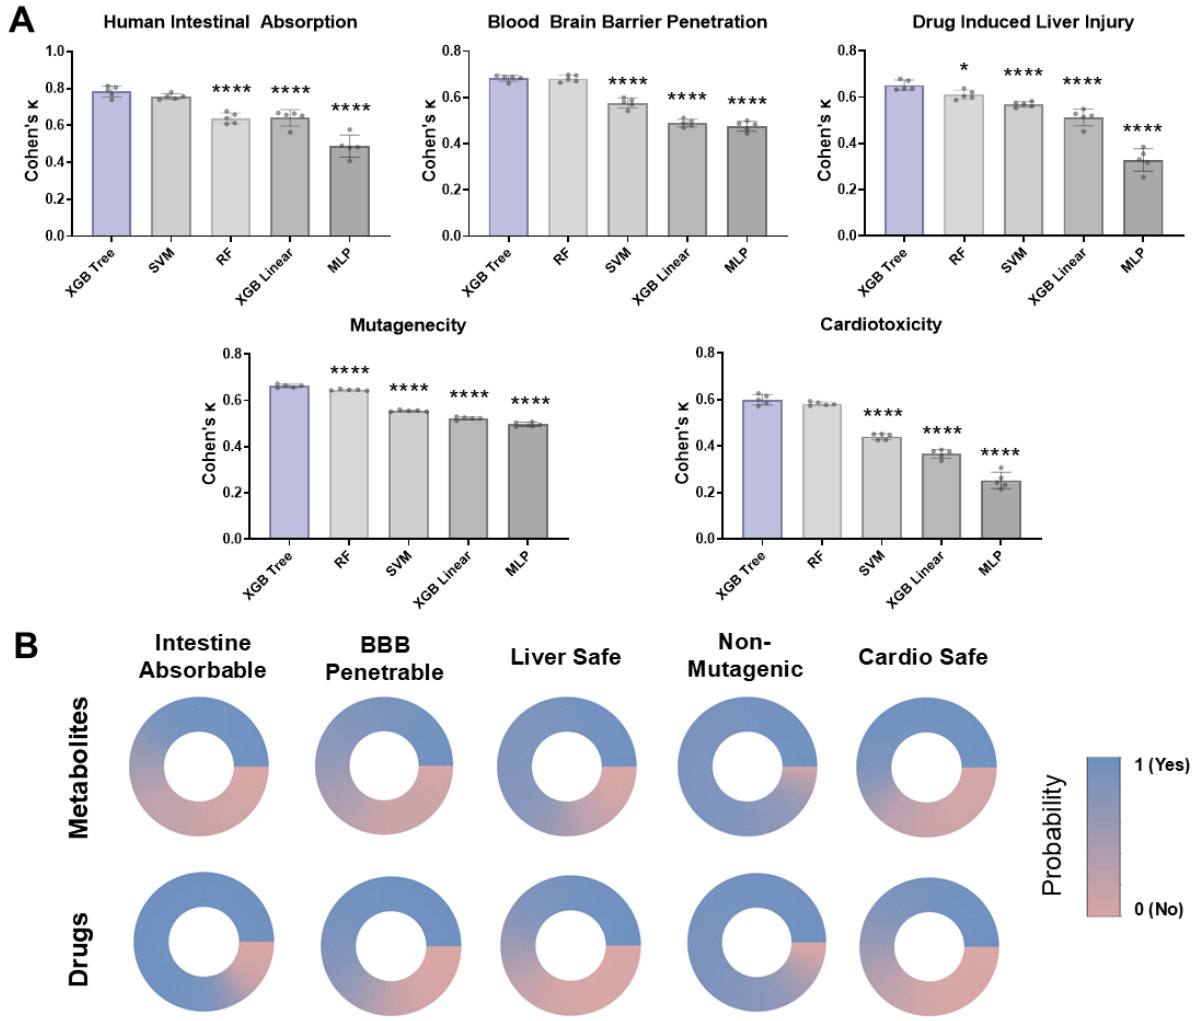

Figure 3. ADMET property prediction of microbiome metabolites. (**A**) Evaluation of different machine learning models. Cohen's κ was used as a performance metric to evaluate five different models using 5 x 10-fold cross-validation, with the XGB tree model consistently showing the best performance. One-way ANOVA followed by Dunnett's post-hoc test was used to assess statistical significance in performance of each model compared to the top performing XGB tree model (* P ≤ 0.05, ** P ≤ 0.01, *** P ≤ 0.001, **** P ≤ 0.0001). (**B**) Predicted probabilities of molecules exhibiting the here investigated properties are represented in a gradient pie chart. Approximately half of the metabolites are predicted to be absorbed through the intestinal barrier. More than half of the metabolites have the potential to penetrate the BBB. Most of the metabolites are predicted to be liver-safe, non-mutagenic, and cardio-safe.

For example, the XGBT model achieved a Cohen's κ of 0.698 for predicting blood-brain barrier (BBB) penetration, which was significantly higher ($P \leq 0.0001$) than that of other models. Its performance was comparable ($P > 0.05$) to the random forest (RF) model in terms of ROC-AUC (XGBT: 0.925; RF: 0.915) and accuracy (XGBT: 0.894; RF: 0.889). Across all endpoints, the XGBT model either performed significantly better than alternative models ($P \leq 0.0001$) in three tasks or demonstrated statistically indistinguishably results in the remaining two, reinforcing its selection as the primary model for prospective prediction tasks. Based on this high retrospective performance of the XGBT model, we hypothesized that these XGBT models could be used to predict ADMET properties of the microbiome metabolites and we set out to analyze predictive trends and validate individual predictions.

Machine learning identifies ADMET trends of microbiome metabolites

We applied the top performing model, XGBT, to predict the modelled ADMET properties of our microbial metabolites (**Fig. 3B**). We first predicted human intestinal absorption from a training dataset containing 906 tested small molecules annotated for their ability to permeate across the human intestine *via* passive diffusion, a criterion crucial for designing orally administered drugs that must bypass the intestinal barrier to reach their intended targets. Similarly, microbiome metabolites that can cross the intestinal barrier may be able to enter systemic circulation and thereby impact host health in tissues beyond the gastrointestinal tract. Several reports have already highlighted instances of microbial metabolites being absorbed across the intestinal barrier[25 26] but it is not yet understood how many and which metabolites might possess the ability to cross this important physiological barrier. Accurate experimental characterization of the intestinal absorption of metabolites requires complex *ex vivo* or *in vivo* experiments and availability of sufficient metabolites for testing, prohibiting the large-scale systematic characterization of the intestinal absorption of microbiome metabolites. Instead, we used our machine learning model to rapidly predict the intestinal absorption propensity of microbiome metabolites. While these predictions might not always be correct, they can prioritize metabolites for further testing and reveal larger overall trends. Our model estimated that about half (50.63%) of microbial metabolites from our database are likely to be absorbed into systemic circulation. Conversely, a substantial proportion (49.37%) of microbiome metabolites might be restricted to the gastrointestinal tract to act locally (**Fig. 3B**). Interestingly, a larger set of the FDA-approved drugs (83.83%) are predicted to be able to permeate across the intestinal membrane, possibly hinting at rational drug design steering this class of molecules towards being more intestinally permeable.

Beyond crossing into systemic circulation through intestinal absorption, another important biological barrier is the Blood-Brain Barrier (BBB). The BBB plays a vital role in protecting the brain from xenobiotics and regulating the passage of substances between the bloodstream and the central nervous system. Some microbiome metabolites have been reported to permeate the BBB and impact regulatory mechanisms in the central nervous system[27], thereby potentially playing a central role in mediating the "gut-brain-axis". The training dataset employed for BBB penetration comprises 1,975 chemical compounds categorized into binary classes based on their ability to permeate the BBB[28 29]. Based on our predictions, the metabolites demonstrate only slightly lower probabilities of penetrating the brain (59.88%) compared to predictions for FDA-approved drugs

(66.9%), highlighting the need for further exploration of the brain-targeting effects of the microbiome metabolites. One example of such a metabolite that has been studied before is trimethylamine-N-oxide (TMAO), a microbiome metabolite that has been reported to cross the BBB and has been associated with various neurodegenerative diseases including Alzheimer's disease and vascular dementia [30]. TMAO was predicted to be BBB-penetrable by our machine learning model, providing another supporting example for the notion that our model can correctly predict the behavior of select metabolites.

Beyond the ability of metabolites to cross biological barriers, we wanted to predict potential toxicities of the microbiome metabolites. For example, the liver is a pivotal organ responsible for regulating xenobiotic metabolism and excretion processes [31]. This makes the liver also vulnerable to damage caused by such xenobiotics. For example, approximately 52.83% of all available drugs have been reported to have caused at least one case of liver injury in patients [32]. This is reflected in our prediction results for the hepatoxic potential of FDA-approved drugs, where 47.33% of the drugs are predicted to be liver-toxic by our machine learning models when using a training dataset of 475 compounds tested for their hepatoxicity[33]. Interestingly, our model suggested that the majority of microbial metabolites (74.2%) are predicted to be liver-safe. This potentially hints at the symbiotic relationship between the microbiome and the host that will steer molecular evolution away from creating microbiome metabolites that could be harmful to their hosts.

Another toxicity we aimed to assess was mutagenicity. To investigate this, we used a training set comprising 7,253 compounds with known mutagenicity profiles[34]. Our machine learning model predicts that the vast majority (95.01%) of microbiome metabolites are unlikely to cause mutagenesis. This result may be somewhat expected, as molecular mutagenicity is commonly assessed using microbial assays such as the Ames test, which commonly relies on assessing mutagenicity in microbes such as *Salmonella* or *E. coli*. Since microbiome metabolites are produced by bacteria, it is plausible that they have evolved to avoid producing compounds that could damage their own genetic material. Nonetheless, given this expected lack of mutagenicity among microbiome metabolites, the low propensity of metabolites predicted to exert mutagenicity by our models further attests to the ability of machine learning models to correctly identify metabolite properties.

Finally, we predicted the expected cardiotoxicity risks of microbiome metabolites by assessing their ability to inhibit the cardiac human Ether-à-go-go related gene (hERG) encoded Kv11.1 potassium channel by relying on publicly-available data for Kv11.1inhibition with 656 entries[35]. The majority of microbiome metabolites were predicted to pose a low risk of hERG inhibition, with 61.85% classified as non-inhibitors of hERG. Similar to the liver toxicity predictions, this could potentially hint at the ability of the microbiome to largely avoid causing damage to their hosts by preventing the creation of cardiotoxic metabolites.

Taken together, these predictions suggested that many microbiome metabolites can be generally considered safe based on predictions of outcomes in three major toxicity assays, but several metabolites are predicted to potentially exert some toxicities that might warrant further attention. Additionally, predictions of biological barrier behavior indicate distinct groups of metabolites, including metabolites that will be restricted to the gastrointestinal tract while others might enter

systemic circulation and even cross into the brain. It is important to note that these trends are based on computational predictions, and that additional validation experiments are necessary to strengthen our confidence in our models and to identify which metabolites indeed exert the predicted properties. Therefore, we set out to perform additional validations and experiments based on these predictions.

Validating barrier-penetration predictions using external datasets

While our machine learning models revealed some interesting trends among microbiome metabolites, it's crucial to further validate the models to ensure that their predictions are accurate. While our retrospective evaluations suggested that the models have acceptable performance (*cf.* **Fig. 3A**), it is possible that these evaluations overestimate performance by testing models on data that at least in-part follows the same underlying distributions as the training data. Instead, evaluations on external data are necessary to assess whether the machine learning models can generalize and make predictions for microbiome metabolites.

For example, by identifying the overlap between metabolites predicted to exhibit a certain property with metabolites that are already known to exhibit this property, we can test the ability of our models to correctly classify positive samples and thereby estimate the model's sensitivity. Although we cannot easily validate whether a model can correctly identify negative samples (specificity) since the absence of a literature report does not necessarily indicate the lack of a property for this metabolite, this overall still allows us to estimate the ability of our models to correctly predict true metabolite properties.

Here, we employed the GUTSY plasma metabolite dataset [36] as an external validation set for our model that classifies the propensity of molecules to undergo human intestinal absorption. The GUTSY online atlas provides an experimentally validated list of metabolites that have been previously detected in human plasma and whose plasma concentration has been shown to correlate with distinct gut microbiome compositions. We hypothesized that a set of microbiome-produced metabolites that can be detected in both the human intestine and human plasma and whose concentration varies with different microbiome compositions is likely to be absorbed through the intestine. Therefore, we matched the metabolites in our database with those detected in the intestine as indicated in the Microbial Metabolites Database (MiMeDB)[37], and then further cross-referenced these metabolites that have been detected in the intestine with the GUTSY online atlas of metabolites detected in plasma. This cross-referencing resulted in an overlapping set of 50 microbiome-derived metabolites that can be detected both in the intestine and in plasma. We then evaluated whether our machine learning model predicted these 50 metabolites to be intestinally absorbed. Our top performing model XGBT demonstrated a high accuracy of 86.0% (**Fig. 4A**). This level of accuracy on the external dataset provides good evidence for the utility of our machine learning model to correctly identify microbiome-derived metabolites that are likely to be absorbed.

In addition, we identified another external validation set for our BBB penetration model, a dataset of metabolites detected in human brains from The National Institute on Aging Genetics of Alzheimer's Disease Data Storage Site (NIAGADS)[38]. Similar to the GUTSY plasma database validation, our hypothesis is that microbial metabolites are likely to be able to penetrate the BBB

if they have been reported to be detectable in the brain. Accordingly, we created a validation set by cross-referencing our microbiome metabolite database with brain metabolites from NIAGADS, resulting in a set of 151 microbiome-derived metabolites that can be found in both datasets. We found that 77.33% of these microbial metabolites that have been detected in the brain are predicted to be BBB-penetrable by our machine learning model, further suggesting that our model accurately predicts such properties for microbiome metabolites (**Fig. 4B**).

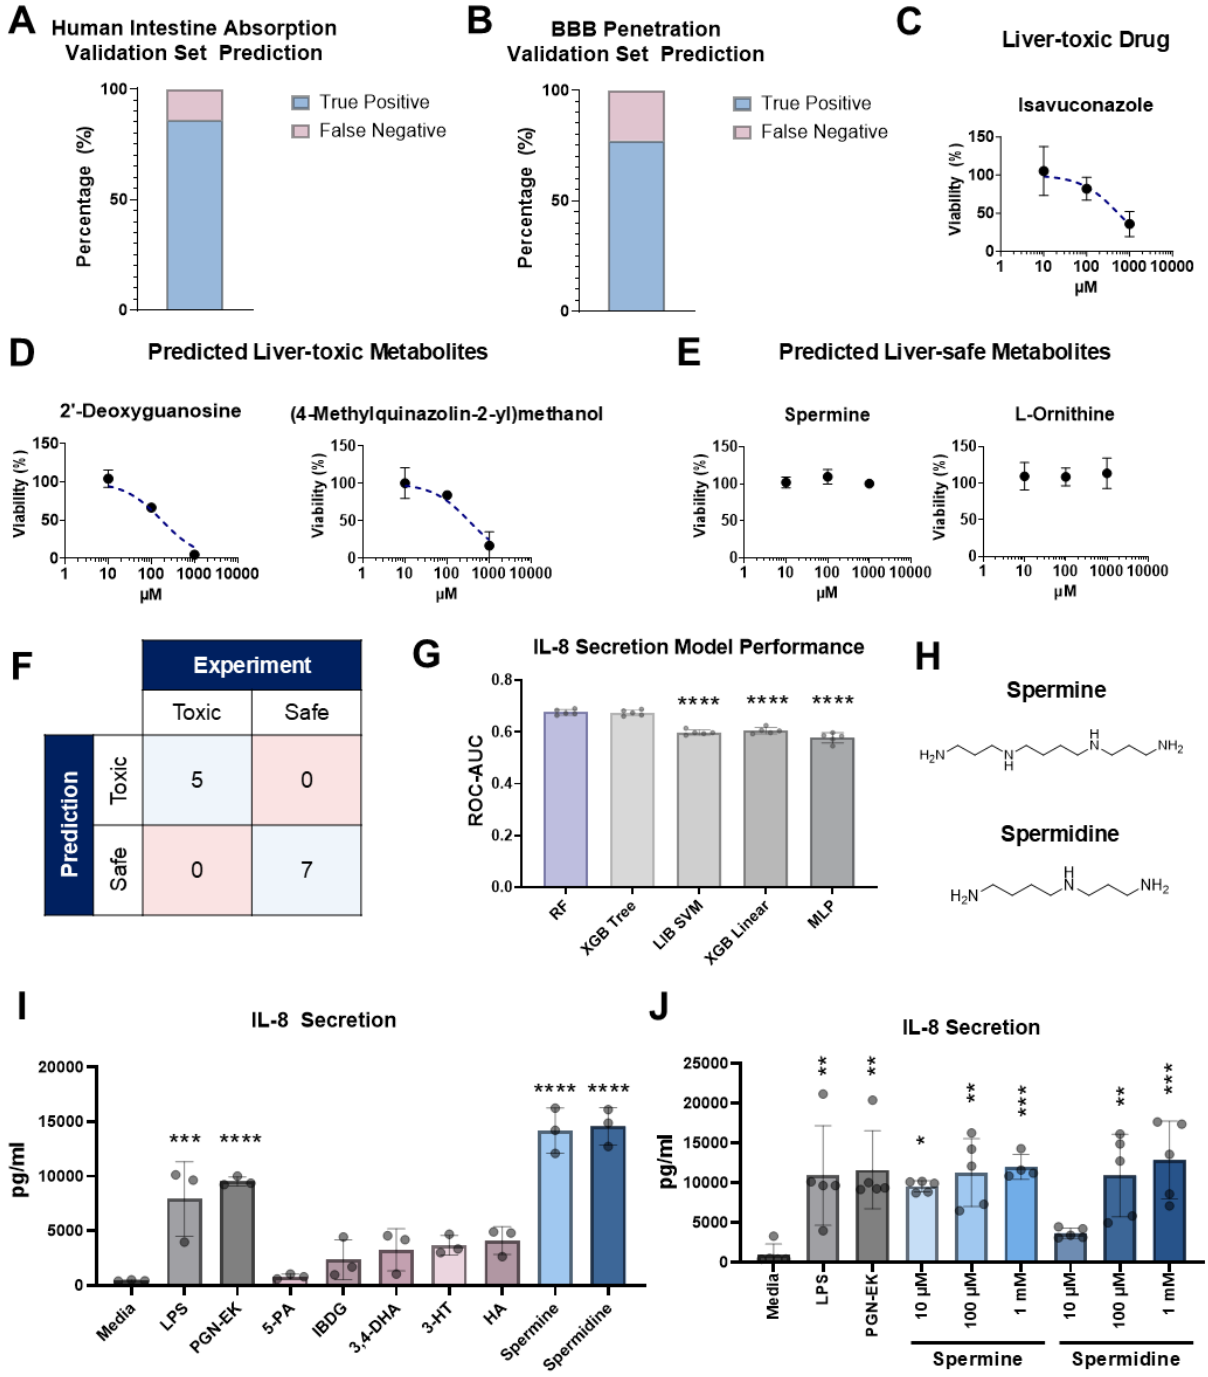

Figure 4. Machine learning model validation using external dataset and in vitro experiments. (**A**) Validation of the human intestinal absorption machine learning model using an external dataset from the GUTSY plasma metabolite database, achieving 86% accuracy. (**B**) Validation of the BBB penetration machine learning model was using the NIAGADS brain metabolite database, demonstrating 77.33% accuracy. (**C-E**) Prospective *in vitro* validation of drug-induced liver injury prediction. HEPG-2 cells were treated with microbial metabolites for 24 hours and cell viability was measured using the MTT assay. The known liver-toxic drug, isavuconazole, served as positive control. Predicted liver-toxic metabolites exhibited dose-dependent, statistically significant cytotoxic effects. Predicted liver-safe metabolites showed no statistically significant cytotoxic effect ($P > 0.05$) at any of the tested concentrations up to 1 mM. An unpaired *t*-test ($P \leq 0.05$) was performed to compare cellular viability at a 1 mM testing concentration between each treatment group and the negative control (media only, no treatment). Each point represents mean values from the experiment (N=3, n=3), and the calculated dose-response curve is shown using a dashed blue line. Error bars indicate standard deviation calculated from biological replicates. (**F**) Confusion matrix for the prospective *in vitro* validation. (**G)** Evaluation of performance of machine learning models predicting IL-8 secretion stimulation. The Random Forest (RF) model outperformed each of the other models (One-way ANOVA followed by Dunnett's post-hoc test; * $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$, **** $P \leq 0.0001$). (**H**) Chemical structures of top predictions spermine and spermidine. (**I**) Results from the IL-8 secretion assay using human PBMC cells, evaluating effect of seven microbiome metabolites at 100 μM (N=1, n=3). Metabolites stimulating IL-8 are shown in orange, and metabolites with no IL-8 secretion are shown in blue. Each point represents an experimental replicate. One-way ANOVA followed by Dunnett's post-hoc test was performed in comparison with the negative control (media only) (* $P \leq 0.05$, ** $P \leq 0.01$, *** $P \leq 0.001$, **** $P \leq 0.0001$). Both spermine and spermidine significantly stimulated IL-8 secretion. (5-PA: 5-Phenylvaleric Acid, IBDG: Indoxyl β-D-Glucoside, 3,4-DHA: 3,4-Dihydroxyhydocinnamic Acid, 3-HT: 3-Hydroxytramine, HA: Hippuric Acid) (**J**) Spermine and spermidine showed a significant dose-dependent effect on IL-8 secretion (N=2, $n \geq 2$) based on the same one-way ANOVA applied as in (I).

In spite of these promising results, it is important to acknowledge some limitations in this analysis. Some metabolites in our database have been reported to be produced both by microbes but also by their hosts. Similarly, some metabolites may not directly cross various biological barriers but instead have primary forms that penetrate and are subsequently metabolized into a specific downstream metabolite *in situ*. In such cases, we are unable to distinguish whether a metabolite is present in a certain tissue because it has trafficked to this tissue from the intestine or because of other mechanisms such as host metabolism. While external datasets offer an economic path to validate predictive models, these inherent limitations might bias at least some of the results. Instead, prospective experiments are a more explicit way to assess model performance and to characterize the properties of specific metabolites. Therefore, we next set out to conduct prospective validation experiments for several of our metabolites.

In vitro validation of drug-induced liver injury predictions

A more rigorous way to evaluate a model than using external test data is prospective experiments, since this enables the validation on novel, unknown data in real-world settings instead of relying upon historical data. We set out to validate the predictions of drug-induced liver injury (DILI) for the metabolites through our in-house *in vitro* cell culture experiments. Utilizing human hepatocyte carcinoma HepG2 cells, a simple but widely employed model for hepatotoxicity assessment[39], we evaluated whether a set of metabolites that are confidently predicted to be hepatotoxic could cause cytotoxicity while metabolites predicted to be safe would have no effect on HepG2 cellular

viability. We curated a list of 12 candidate metabolites (**Fig. S2, Table S3**) based on our machine learning model prediction for prospective experimental evaluation using high predictive confidence while ensuring commercial availability and sufficient aqueous solubility of the metabolites to enable experimental testing. We purchased seven metabolites predicted to be liver-safe and five metabolites predicted to be liver-toxic by our model for testing. Specifically, we prioritized metabolites with predicted probabilities of being liver-safe either well above 0.9 or below 0.15 to enable evaluation of both predicted safe and potentially unsafe compounds. One metabolite, Indoxyl β-D-Glucoside with a predicted probability of 0.487, was included as a borderline case to further probe the model's performance.

All of these selected metabolites were then tested for their cytotoxicity on HepG2 cells at three different molecular concentrations. Cells were treated for 24 hours, and we assessed cell viability using the MTT assay. Remarkably, all six metabolites predicted to be toxic demonstrated dose-dependent liver cell toxicity (**Fig. 4D, Fig. S2A, Table S3**), exhibiting a similar dose-dependent pattern compared to the well-known liver-toxic drug isavuconazole which was used as a positive control in our experiments (**Fig. 4C**). Conversely, none of the metabolites predicted to be safe showed any statistically significant cytotoxicity at any of the tested concentrations (**Fig. 4E, Fig. S2B, Table S3**).

Some of our predicted liver-toxic metabolites have previously published evidence of hepatotoxicity. For example, phloretin has been shown to be liver toxic in rats, as noted in recent research [40]. Although we did not find direct evidence of liver toxicity for 2'-deoxyguanosine in the literature, oxidatively modified form, 8-hydroxy-2'-deoxyguanosine, is known to exhibit liver toxicity [41] [42]. In contrast, to the best of our knowledge, levomefolic acid and xanthosine have no prior reported evidence of associated liver toxicity but nevertheless showed dose-dependent cytotoxicity in our assays in full agreement with our predictions. Although we were only able to test 12 metabolites in a simple *in vitro* model of liver cell toxicity, the achieved 100% accuracy suggests our model is robust for further characterization and to reveal trends among metabolites (**Fig. 4F**). The novel associations such as the *in vitro* effects of levomefolic acid and xanthosine on HepG2 viability add to the body of work for these metabolites and warrants further research.

Predicting IL-8 stimulatory effects of metabolites

Having studied and further validated our machine learning models' prediction of ADMET properties of metabolites, we aimed to next address an immunological question using our platform. It is known that certain microbial metabolites can influence the host immune system to provide important health effects, such as microbiome-derived inosine's potential role to support the response of mice to immunotherapy[43]. However, given the complexity of the immune system and the microbiome, our understanding of the interplay between microbiome metabolites and the immune system continues to be sparse.

To create a case study to assess the potential of machine learning to support the identification of immunomodulatory microbiome metabolites, we built a machine learning model to predict whether a small molecule could induce interleukin 8 (IL-8) secretion using a training set comprising 1321 data points sourced from the publicly available PubChem repository. The training

set includes binary labels denoting labels for small molecules to be active (stimulating IL-8 secretion) and inactive (not stimulating IL-8 secretion).

Similar to our previous model selection strategy, we trained and evaluated multiple machine learning model architectures on this data. We found that RF and XGBT outperformed ($P < 0.001$) all other model architectures on this data in cross-validation experiments according to the ROC-AUC and Cohen's κ metrics (**Fig 4G, S3**) while RF had consistently had the highest absolute value for these metrics across all models. Subsequently, we applied this RF model to our microbiome metabolite database to predict new associations for further testing. To validate these predictions, we selected three positively predicted metabolites as well as four negatively predicted metabolites for further *in vitro* experiments (**Fig. 4I**, **Table S4**). We assessed IL-8 stimulation by treating human peripheral blood mononuclear cells (PBMCs) with each of the seven metabolites at a concentration of 100 μM for 24 hours (**Fig. 4I**). Experimental validation demonstrated high predictive accuracy (85.71%) with two true positives - spermidine and spermine - showing IL-8 stimulation. Both metabolites are well-known for their anti-inflammatory properties, [44] [45] but the model correctly predicted that they could stimulate IL-8 secretion with high confidence (probability for spermidine = 0.76, probability for spermine = 0.64; **Fig. 4H, Table S4**). None of the negative predictions showed any IL-8 stimulation and only one false positive, Indoxyl β-D-Glucoside, was identified. This showcases the potential of machine learning models to identify immunomodulatory effects of microbiome metabolites, with large potential to accelerate future microbiome research.

We subsequently performed a dose-response assay at three different concentrations of spermine and spermidine and found that their IL-8 release stimulation changed in a concentration-dependent manner (**Fig. 4J**). To further contextualize the IL-8 secretion stimulation by spermine and spermidine, we tested these metabolites for their potential to stimulate the secretion of other cytokines, including IL-1α, IL-6, IL-12p70, and TNFα. No significant stimulation was observed for any of the tested cytokines, indicating a striking selectivity of spermine's and spermidines' immunostimulatory potential.

Taken together, this case-study highlights the potential of machine learning to facilitate the study of the interactions between the microbiome and the immune system. Additionally, the previously unknown, selective IL-8 release stimulation of the two well-studied anti-inflammatory molecules spermine and spermidine warrants further attention with potentially important implications for their further development as biomarkers or supplements.

CONCLUSION

The microbiome is increasingly recognized as a crucial component of host physiology, impacting various processes associated with human health and disease. In particular, microbiome metabolites are receiving increasing attention due to their potential role in mediating these interactions. However, only a minute fraction of microbiome metabolites has been characterized so far. Since the study of the biological effects of metabolites is laborious and expensive, we here proposed, prototyped, and evaluated machine learning models that can triage experiments to accelerate the identification of biological effects of microbiome metabolites.

Specifically, we developed several machine learning models capable of predicting a wide range of properties of human microbiome metabolites, ranging from ADMET properties such as intestinal absorption and hepatoxicity to immunological properties exemplified by IL-8 secretion stimulation. The models were validated through a series of external datasets and in-house experiments, demonstrating strong predictive performance and robustness. Notably, the DILI prediction model achieved 100% accuracy, while the IL-8 secretion model reached 85.71%, resulting in an overall accuracy of 92.86% across prospective validations. These results suggest that machine learning algorithms are well suited to anticipate microbiome metabolite properties, although occasional false positives may still occur. Furthermore, our work has already led to the discovery of several previously unknown properties of multiple microbiome metabolites - including the selective stimulation of IL-8 secretion through spermine and spermidine.

The quality of such platforms depends on the quality of the machine learning models, training data for specific properties, and the coverage of the microbiome metabolite database. With increasing focus of the research community on artificial intelligence, high-throughput experimental platforms for large data generation, and microbiome metabolomics, we expect the utility and quality of such machine learning-driven predictive platforms that can guide the study of microbiome metabolites and their effects to continuously improve. Furthermore, many of the datasets used in machine learning are derived from *in vitro* experiments, which might not always correlate with *in vivo* outcomes of metabolites. Nevertheless, the prioritization of metabolites for *in vitro* and subsequent *in vivo* experiments could constitute a rapid alternative to untargeted screening experiments.

In the future, we expect the machine learning-assisted functional annotation of microbiome metabolites to facilitate the identification of biomarkers, which can aid in patient diagnostics and prognostics as well as triage patient populations to better serve patients with distinct microbiome composition. Ultimately, we hope that such work can also enable the development of potential therapeutic interventions derived from microbiome metabolites to borrow from our microbiome's symbiotic molecular arsenal to derive novel therapies and adjuvants to treat and prevent diseases.

METHODS

Experimental model details

Cell lines

All cells were cultured at 37 ºC in a 5% CO_2 atmosphere. HEPG2 cells were cultured using DMEM medium supplemented with 10% fetal bovine serum, 100 U/mL penicillin, and 100 µg/mL streptomycin. Two primary human peripheral blood mononuclear cells (PBMCs) were used. The first sample (PBMC060822AH) was obtained from a male Caucasian (46 years old). The second sample was obtained from ATCC (PCS-800-011, Lot 8032322), from a female Caucasian (72 years old). PBMCs were maintained using RMPI-1640 based Lymph-1 lymphocyte medium (Labscoop Lymph-1) supplemented with 10% fetal bovine serum, 100 U/mL penicillin, and 100 µg/mL streptomycin.

Method details

Curation of the microbiome metabolite database

We curated a database of microbiome metabolites from five different publicly available sources, with a focus on selecting compounds of microbiome origin or human–microbiome co-metabolites that are chemically modified by commensal microbes. The databases are: Human Metabolome Database (HMDB)[46], Virtual Metabolic Human (VMH)[47], KNApsack Comprehensive Species-Metabolite Relationship Database (KNApSAcK) [48], Microbial Metabolites Database (MiMeDB)[49], and the Metabolomics Data Explorer[50]. From HMDB and VMH, only metabolites with the source annotated as “bacteria” were included. In KNApSAcK, we filtered for metabolites reported to originate from both human and bacterial sources. From MiMeDB, we selected metabolites categorized as either “primary” (originating solely from microbial biosynthetic pathways) or “co-metabolites” (produced or modified by both human and microbiome metabolic processes). All metabolites obtained from the Metabolomics Data Explorer were included based on their detection in *in vitro* culture of commensal microbes. Each metabolite data entry was annotated with the microbial strains that are known to produce this metabolite along with additional identifiers such as their chemical name, CAS number, and InChI key, if this information was available in any of the data sources. Data that contained missing values and duplicates were removed. The final database captures the chemical structures (stored as canonical Simplified Molecular Input Line Entry System; SMILES) of 2,645 metabolites.

Datasets for machine learning model training

The dataset of FDA-approved drugs was sourced from a review paper[51], originally derived from DrugBank 2020[52] . Training datasets for ADMET property prediction were obtained from the Therapeutics Data Commons (TDCommons)[53]. Additionally, the training set for IL-8 secretion was obtained from PubChem (Assay IDs 651758, 327254, 598346)[54] [55] [56] while removing compounds labeled as “Inconclusive” and “Inhibitor” and labeling "Active" compounds as the

positive class and “Inactive" compounds as negative class for IL-8 secretion. For all the datasets, data entries containing missing values were removed and duplicates were retained only when they were consistently annotated according to their labels.

Analysis of “drug-like” properties

Chemical properties (LogP, molecular weight, hydrogen bond donors, hydrogen bond acceptors, topological surface area, rotatable bonds) were calculated using RDKit nodes (Version 2023.09.2) in KNIME (Version 5.2.3.). The t-SNE plot was generated using Morgan fingerprints (1024 bit, radius of 2) and RDKit property descriptors. Kernel density estimate (KDE) plots were generated using the Seaborn library in Python.

Designing and cross-validation of machine learning models

All the classification machine learning models were developed in the KNIME data analysis software (KNIME 5.2.3). Five machine learning models were tested for each property prediction: random forest (RF), XGBoost linear, XGBoost tree, multi-layer perceptron (MLP), and Support Vector Machine (SVM). All models were used with default parameters. Molecules were described with Morgan fingerprints (1024 bit, radius of 2) and RDKit descriptors. Descriptors were standardized before model training on the training data and the same standardization was applied to the test data. Models were evaluated through stratified 5 X 10-fold cross-validation evaluations for model selection. Performance was quantified using the metrics ROC-AUC, accuracy, and Cohen’s κ.

Cytotoxicity assay of Hep G2 cells

Hep G2 cells were plated in a 96-well plate at a density of 300,000 cells per well and allowed to adhere overnight at 37°C in 5% CO_2. The following day, the cells were treated with metabolites or positive control cytotoxic drug at three different testing concentrations (1 mM, 100 µM, 10 µM) in culture medium, followed by a 24-hour incubation period. After incubation, MTT solution was added to each well and incubated for 4 hours at 37°C in 5% CO_2. The medium was then removed, and 100 µL of DMSO was added to each well. After shaking the plate at 400 rpm for 10 minutes at room temperature, the absorbance at 570 nm was measured using a plate reader (Tecan Infinite M Plex).

Cytokine stimulation assay

Human PBMCs were plated in a 96-well plate at a density of 150,000 cells per well and allowed to adhere overnight at 37°C in 5% CO_2. The following day, cells were treated with metabolites at different concentrations (1 mM, 100 µM, 10 µM). 1 µg/mL LPS (Invivogen tlrl-eblps) and 1 µg/mL PGN-EK (Invivogen tlrl-pgnek) served as positive controls. All treatments were dissolved in

culture medium. Cells were incubated with the treatments for a 24-hour period. After 24 hours, the media was removed and stored at -80°C until analysis. The IL-1α, IL-6, IL-8, IL-12p70, and TNFα levels were measured using a Milliplex analyte kit (EMD Millipore). Cytokine stimulation experiments were performed in collaboration with the UNC CGIBD Advanced Analytics core facility and ZenBio Inc. in fee-for-service agreements.

Statistical analysis

All data were analyzed with the GraphPad Prism software. Machine learning performance was evaluated statistically using a two-way or a one-way ANOVA followed by Dunnett's or Tukey's post-hoc tests. Cell assay results were evaluated using a one-way ANOVA with Dunnett's post-hoc test or an unpaired T-test. Error bars denote standard deviation across biological replicates in all figures.

Acknowledgements

This work was supported, in part, by a Flash Grant from NC Biotech (2021-FLG-3819), a UNC CGIBD Pilot Award (NIH NIDDK DK034987), a Duke Cancer Institute and Duke Microbiome Center Pilot Award (NIH NCI CA014236), and the Engineering Research Center for Precision Microbiome Engineering (NSF EEC-2133504). Cytokine release studies were conducted in collaboration with the UNC CGIBD Advanced Analytics core facility and ZenBio Inc. in a fee-for-service agreement. Z.F. is supported by the Department of Defense (DoD) through the National Defense Science and Engineering Graduate (NDSEG) Fellowship Program.

Competing Interests

D.R. acts as a consultant to the pharmaceutical and biotechnology industry, as a mentor for Start2, and on the scientific advisory board of Areteia Therapeutics.

REFERENCES

[1] Sharon, G., Garg, N., Debelius, J., Knight, R., Dorrestein, P.C., and Mazmanian, S.K. (2014). Specialized metabolites from the microbiome in health and disease. *Cell Metab.* 20, 719–730. https://doi.org/10.1016/j.cmet.2014.10.016.

[2] Li, Z., Quan, G., Jiang, X., Yang, Y., Ding, X., Zhang, D., Wang, X., Hardwidge, P.R., Ren, W., and Zhu, G. (2018). Effects of metabolites derived from gut microbiota and hosts on pathogens. *Front. Cell. Infect. Microbiol.* 8, 314. https://doi.org/10.3389/fcimb.2018.00314.

[3] Afzaal, M., Saeed, F., Shah, Y.A., Ahmad, N., Imran, M., Hussain, M., Tufail, T., and Bader-Ul-Ain, H. (2022). Human gut microbiota in health and disease: Unveiling the relationship. *Front. Microbiol.* 13, 999001. https://doi.org/10.3389/fmicb.2022.999001.

[4] Krishnan, S., Ding, Y., Saedi, N., Choi, M., Sridharan, G.V., Sherr, D.H., Yarmush, M.L., and Upreti, M. (2018). Gut microbiota–derived tryptophan metabolites modulate inflammatory response in hepatocytes and macrophages. *Cell Rep.* 23, 1099–1111. https://doi.org/10.1016/j.celrep.2018.03.110.

[5] Mager, L.F., Burkhard, R., Pett, N., Cooke, N.C.A., Brown, K., Ramay, H., Paik, S., Stagg, J., Groves, R.A., Gallo, M., et al. (2020). Microbiome-derived inosine modulates response to checkpoint inhibitor immunotherapy. *Science* 369, 1481–1489. https://doi.org/10.1126/science.abc3421.

[6] Heianza, Y., Ma, W., Manson, J.E., Rexrode, K.M., and Qi, L. (2017). Gut microbiota metabolites and risk of major adverse cardiovascular disease events and death: A systematic review and meta-analysis of prospective studies. *J. Am. Heart Assoc.* 6, e004947. https://doi.org/10.1161/JAHA.116.004947.

[7] Gold, A., and Zhu, J. (2022). Not just a gut feeling: A deep exploration of functional bacterial metabolites that can modulate host health. *Gut Microbes* 14, 2125734. https://doi.org/10.1080/19490976.2022.2125734.

[8] Nicolas, G.R., and Chang, P.V. (2019). Deciphering the chemical lexicon of host–gut microbiota interactions. *Trends Pharmacol. Sci.* 40, 430–445. https://doi.org/10.1016/j.tips.2019.04.006.

[9] Mullowney, M.W., Duncan, K.R., Elsayed, S.S., Shoichet, B.K., and Carlson, E.E. (2023). Artificial intelligence for natural product drug discovery. *Nat. Rev. Drug Discov.* 22, 895–916. https://doi.org/10.1038/s41573-023-00774-7.

[10] Rodrigues, T., Reker, D., Schneider, P., and Schneider, G. (2016). Counting on natural products for drug design. *Nat. Chem.* 8, 531–541. https://doi.org/10.1038/nchem.2479.

[11] Mullowney, M.W., Duncan, K.R., Elsayed, S.S., Shoichet, B.K., and Carlson, E.E. (2023). Artificial intelligence for natural product drug discovery. *Nat. Rev. Drug Discov.* 22, 895–916. https://doi.org/10.1038/s41573-023-00774-7.

[12] Wishart, D.S., Feunang, Y.D., Marcu, A., Guo, A.C., Liang, K., Vázquez-Fresno, R., Sajed, T., Johnson, D., Li, C., Karu, N., et al. (2018). HMDB 4.0: The human metabolome database for 2018. *Nucleic Acids Res.* 46, D608–D617. https://doi.org/10.1093/nar/gkx1089.

[13] Brunk, E., Sahoo, S., Zielinski, D.C., Altunkaya, A., Dräger, A., Mih, N., Gatto, F., Nilsson, A., Preciat Gonzalez, G.A., Aurich, M.K., et al. (2018). Recon3D enables a three-dimensional view of gene variation in human metabolism. *Nat. Biotechnol.* 36, 272–281. https://doi.org/10.1038/nbt.4072.

[14] Nakamura, Y., Afendi, F.M., Parvin, A.K., Ono, N., Tanaka, K., Morita, A.H., Sato, T., Sugiura, T., Altaf-Ul-Amin, M., and Kanaya, S. (2014). KNApSAcK metabolite activity database for retrieving the relationships between metabolites and biological activities. *Plant Cell Physiol.* 55, e7. https://doi.org/10.1093/pcp/pct178.
[15] Wishart, D.S., Oler, E., Peters, H., Guo, A.C., Girod, S., Han, S., Saha, S., Lui, V.W.Y., LeVatte, M., Gautam, V., et al. (2023). MiMeDB: The human microbial metabolome database. *Nucleic Acids Res.* 51, D611–D620. https://doi.org/10.1093/nar/gkac868.
[16] Han, S., Van Treuren, W., Fischer, C.R., Merrill, B.D., DeFelice, B.C., Sanchez, J.M., Higginbottom, S.K., Guthrie, L., Fall, L.A., Dodd, D., et al. (2021). A metabolomics pipeline for the mechanistic interrogation of the gut microbiome. *Nature* 595, 415–420. https://doi.org/10.1038/s41586-021-03707-9.
[17] Hamamah, S., Aghazarian, A., Nazaryan, A., Hajnal, A., and Covasa, M. (2022). Role of microbiota–gut–brain axis in regulating dopaminergic signaling. *Biomedicines* 10, 436. https://doi.org/10.3390/biomedicines10020436.
[18] Mullowney, M.W., Duncan, K.R., Elsayed, S.S., Shoichet, B.K., and Carlson, E.E. (2023). Artificial intelligence for natural product drug discovery. *Nat. Rev. Drug Discov.* 22, 895–916. https://doi.org/10.1038/s41573-023-00774-7.
[19] Huang, W., Brewer, L.K., Jones, J.W., Nguyen, A.T., Marcu, A., Wishart, D.S., Oglesby-Sherrouse, A.G., Kane, M.A., and Wilks, A. (2018). PAMDB: A comprehensive Pseudomonas aeruginosa metabolome database. *Nucleic Acids Res.* 46, D575–D580. https://doi.org/10.1093/nar/gkx1061.
[20] Wilson, M.G., and Pandey, S. (2024). Pseudomonas aeruginosa. In: StatPearls [Internet]. Treasure Island (FL): StatPearls Publishing. https://www.ncbi.nlm.nih.gov/books/NBK557831/.
[21] Lipinski, C.A., Lombardo, F., Dominy, B.W., and Feeney, P.J. (2001). Experimental and computational approaches to estimate solubility and permeability in drug discovery and development settings. *Adv. Drug Deliv. Rev.* 46, 3–26. https://doi.org/10.1016/S0169-409X(00)00129-0.
[22] Ekins, S., Waller, C.L., Swaan, P.W., Cruciani, G., Wrighton, S.A., and Wikel, J.H. (2000). Progress in predicting human ADME parameters in silico. *J. Pharmacol. Toxicol. Methods* 44, 251–272. https://doi.org/10.1016/s1056-8719(00)00109-x.
[23] Johnson, A.M., and Maggiora, G.M. (1990). *Concepts and Applications of Molecular Similarity.* New York: John Wiley & Sons. https://doi.org/10.1002/9781118662563.
[24] TDCommons (2023). TDCommons: Drug-like molecule dataset. *Therapeutics Data Commons*. https://doi.org/10.7910/DVN/21LKWG.
[25] Uchimura, Y., Fuhrer, T., Li, H., Lawson, M.A., Zimmermann, M., Yilmaz, B., Zindel, J., Ronchi, F., Sorribas, M., Hapfelmeier, S., et al. (2018). Antibodies set boundaries limiting microbial metabolite penetration and the resultant mammalian host response. Immunity 49, 545–559.e5. https://doi.org/10.1016/j.immuni.2018.08.008.
[26] Iyer, N., and Corr, S.C. (2021). Gut microbial metabolite-mediated regulation of the intestinal barrier in the pathogenesis of inflammatory bowel disease. *Nutrients* 13, 4259. https://doi.org/10.3390/nu13124259.
[27] Swer, N.M., Venkidesh, B., Murali, T.S., and Prasad, M.N.V. (2023). Gut microbiota-derived metabolites and their importance in neurological disorders. *Mol. Biol. Rep.* 50, 1663–1675. https://doi.org/10.1007/s11033-023-08272-5.

[28] Martins, I.F., Teixeira, A.L., Pinheiro, L., and Falcão, A.O. (2012). A Bayesian approach to in silico blood–brain barrier penetration modeling. *J. Chem. Inf. Model.* 52, 1686–1697. https://doi.org/10.1021/ci300124c.
[29] Wu, Z., Ramsundar, B., Feinberg, E.N., Gomes, J., Geniesse, C., Pappu, A.S., Leswing, K., and Pande, V. (2018). MoleculeNet: A benchmark for molecular machine learning. *Chem. Sci.* 9, 513–530. https://doi.org/10.1039/C7SC02664A.
[30] Parker, A., Fonseca, S., and Carding, S.R. (2020). Gut microbes and metabolites as modulators of blood–brain barrier integrity and brain health. *Gut Microbes* 11, 135–157. https://doi.org/10.1080/19490976.2019.1638722.
[31] Andrade, R.J., Chalasani, N., Björnsson, E.S., Suzuki, A., Kullak-Ublick, G.A., and Watkins, P.B. (2019). Drug-induced liver injury. *Nat. Rev. Dis. Primers* 5, 58. https://doi.org/10.1038/s41572-019-0105-0.
[32] Björnsson, E.S. (2016). Hepatotoxicity by drugs: The most common implicated agents. *Int. J. Mol. Sci.* 17, 224. https://doi.org/10.3390/ijms17020224.
[33] Xu, Y., Dai, Z., Chen, F., Gao, S., Pei, J., and Lai, L. (2015). Deep learning for drug-induced liver injury. *J. Chem. Inf. Model.* 55, 2085–2093. https://doi.org/10.1021/acs.jcim.5b00238.
[34] Xu, C., Cheng, F., Chen, L., Du, Z., Li, W., Liu, G., Lee, P.W., and Tang, Y. (2012). In silico prediction of chemical Ames mutagenicity. *J. Chem. Inf. Model.* 52, 2840–2847. https://doi.org/10.1021/ci300400a.
[35] Du, F., Yu, H., Zou, B., Babcock, J., Long, S., and Li, M. (2011). hERGCentral: A large database to store, retrieve, and analyze compound–human Ether-à-go-go related gene channel interactions to facilitate cardiotoxicity assessment in drug development. *Assay Drug Dev. Technol.* 9, 580–588. https://doi.org/10.1089/adt.2011.0425.
[36] Dekkers, K.F., Sayols-Baixeras, S., Baldanzi, G., Nowak, C., Hammar, U., Nguyen, D., Varotsis, G., Brunkwall, L., Nielsen, N., Eklund, A.C., et al. (2022). An online atlas of human plasma metabolite signatures of gut microbiome composition. *Nat. Commun.* 13, 5370. https://doi.org/10.1038/s41467-022-33050-0.
[37] Wishart, D.S., Oler, E., Peters, H., Guo, A., Girod, S., Han, S., Saha, S., Lui, V.W., LeVatte, M., Gautam, V., et al. (2023). MiMeDB: The human microbial metabolome database. *Nucleic Acids Res.* 51, D611–D620. https://doi.org/10.1093/nar/gkac868.
[38] NIAGADS (2023). NG00131 – Large scale genetics study on CSF and brain metabolite levels. *National Institute on Aging Genetics of Alzheimer's Disease Data Storage Site (NIAGADS).* https://doi.org/10.60859/cbxy-g768.
[39] Ramirez, T., Strigun, A., Verlohner, A., Huener, H.A., Peter, E., Herold, M., Bordag, N., Mellert, W., Walk, T., Spitzer, M., et al. (2018). Prediction of liver toxicity and mode of action using metabolomics in vitro in HepG2 cells. *Arch. Toxicol.* 92, 893–906. https://doi.org/10.1007/s00204-017-2079-6.
[40] Itou da Silva, F.S., Veiga Bizerra, P.F., Mito, M.S., Constantin, R.P., Klosowski, E.M., Lima de Souza, B.T., Moreira da Costa Menezes, P.V., Alves Bueno, P.S., Nanami, L.F., Marchiosi, R., et al. (2022). The metabolic and toxic acute effects of phloretin in the rat liver. *Chem. Biol. Interact.* 364, 110054. https://doi.org/10.1016/j.cbi.2022.110054.
[41] Ichiba, M., Maeta, Y., Mukoyama, T., Saeki, T., Yasui, S., Kanbe, T., Okano, J., Tanabe, Y., Hirooka, Y., Yamada, S., et al. (2003). Expression of 8-hydroxy-2'-deoxyguanosine in chronic liver disease and hepatocellular carcinoma. *Liver Int.* 23, 338–345. https://doi.org/10.1034/j.1478-3231.2003.00868.x.

[42] Hahm, J.Y., Park, J., Jang, E.S., Kim, H.S., and Kim, J. (2022). 8-Oxoguanine: From oxidative damage to epigenetic and epitranscriptional modification. *Exp. Mol. Med.* 54, 1626–1642. https://doi.org/10.1038/s12276-022-00822-z.
[43] Mager, L.F., Burkhard, R., Pett, N., Cooke, N.C.A., Brown, K., Ramay, H., Paik, S., Stagg, J., Groves, R.A., Gallo, M., et al. (2020). Microbiome-derived inosine modulates response to checkpoint inhibitor immunotherapy. *Science* 369, 1481–1489. https://doi.org/10.1126/science.abc3421.
[44] Zhang, M., Caragine, T., Wang, H., Cohen, P.S., Botchkina, G., Soda, K., Bianchi, M., Ulrich, P., Cerami, A., Sherry, B., and Tracey, K.J. (1997). Spermine inhibits proinflammatory cytokine synthesis in human mononuclear cells: A counterregulatory mechanism that restrains the immune response. *J. Exp. Med.* 185, 1759–1768. https://doi.org/10.1084/jem.185.10.1759.
[45] Choi, Y.H., and Park, H.Y. (2012). Anti-inflammatory effects of spermidine in lipopolysaccharide-stimulated BV2 microglial cells. *J. Biomed. Sci.* 19, 31. https://doi.org/10.1186/1423-0127-19-31.
[46] Wishart, D.S., Feunang, Y.D., Marcu, A., Guo, A.C., Liang, K., Vázquez-Fresno, R., Sajed, T., Johnson, D., Li, C., Karu, N., et al. (2018). HMDB 4.0: The human metabolome database for 2018. *Nucleic Acids Res.* 46, D608–D617. https://doi.org/10.1093/nar/gkx1089.
[47] Brunk, E., Sahoo, S., Zielinski, D.C., Altunkaya, A., Dräger, A., Mih, N., Gatto, F., Nilsson, A., Preciat Gonzalez, G.A., Aurich, M.K., et al. (2018). Recon3D enables a three-dimensional view of gene variation in human metabolism. *Nat. Biotechnol.* 36, 272–281. https://doi.org/10.1038/nbt.4072.
[48] Nakamura, Y., Afendi, F.M., Parvin, A.K., Ono, N., Tanaka, K., Morita, A.H., Sato, T., Sugiura, T., Altaf-Ul-Amin, M., and Kanaya, S. (2014). KNApSAcK metabolite activity database for retrieving the relationships between metabolites and biological activities. *Plant Cell Physiol.* 55, e7. https://doi.org/10.1093/pcp/pct178.
[49] Wishart, D.S., Oler, E., Peters, H., Guo, A., Girod, S., Han, S., Saha, S., Lui, V.W., LeVatte, M., Gautam, V., et al. (2023). MiMeDB: The human microbial metabolome database. *Nucleic Acids Res.* 51, D611–D620. https://doi.org/10.1093/nar/gkac868.
[50] Han, S., Van Treuren, W., Fischer, C.R., Merrill, B.D., DeFelice, B.C., Sanchez, J.M., Higginbottom, S.K., Guthrie, L., Fall, L.A., Dodd, D., et al. (2021). A metabolomics pipeline for the mechanistic interrogation of the gut microbiome. *Nature* 595, 415–420. https://doi.org/10.1038/s41586-021-03707-9.
[51] Fralish, Z., Chen, A., Khan, S., Lee, J., Ho, J., and Shoichet, B.K. (2024). The landscape of small-molecule prodrugs. *Nat. Rev. Drug Discov.* 23, 365–380. https://doi.org/10.1038/s41573-024-00914-7.
[52] Wishart, D.S., Feunang, Y.D., Guo, A.C., Lo, E.J., Marcu, A., Grant, J.R., Sajed, T., Johnson, D., Li, C., Sayeeda, Z., et al. (2018). DrugBank 5.0: A major update to the DrugBank database for 2018. *Nucleic Acids Res.* 46, D1074–D1082. https://doi.org/10.1093/nar/gkx1037.
[53] TDCommons (2023). TDCommons: Drug-like molecule dataset. *Therapeutics Data Commons*. https://doi.org/10.7910/DVN/21LKWG.
[54] National Center for Biotechnology Information (2024). PubChem Bioassay Record for AID 651758, qHTS assay to identify small molecules that stimulate interleukin-8 (IL-8) secretion. Source: National Center for Advancing Translational Sciences (NCATS). Retrieved September 29, 2024 from https://pubchem.ncbi.nlm.nih.gov/bioassay/651758.

[55] National Center for Biotechnology Information (2025). PubChem Bioassay Record for AID 327254, Source: ChEMBL. Retrieved February 16, 2025 from https://pubchem.ncbi.nlm.nih.gov/bioassay/327254.
[56] National Center for Biotechnology Information (2025). PubChem Bioassay Record for AID 598346, Source: ChEMBL. Retrieved February 16, 2025 from https://pubchem.ncbi.nlm.nih.gov/bioassay/598346.