

<https://doi.org/10.1038/s42003-024-06797-z>

# Semi-supervised meta-learning elucidates understudied molecular interactions

Check for updates

You Wu<sup>1</sup>, Li Xie<sup>2</sup>, Yang Liu<sup>2</sup> & Lei Xie<sup>1,2,3</sup> ✉

Many biological problems are understudied due to experimental limitations and human biases. Although deep learning is promising in accelerating scientific discovery, its power compromises when applied to problems with scarcely labeled data and data distribution shifts. We develop a deep learning framework—Meta Model Agnostic Pseudo Label Learning (MMAPLE)—to address these challenges by effectively exploring out-of-distribution (OOD) unlabeled data when conventional transfer learning fails. The uniqueness of MMAPLE is to integrate the concept of meta-learning, transfer learning and semi-supervised learning into a unified framework. The power of MMAPLE is demonstrated in three applications in an OOD setting where chemicals or proteins in unseen data are dramatically different from those in training data: predicting drug-target interactions, hidden human metabolite-enzyme interactions, and understudied interspecies microbiome metabolite-human receptor interactions. MMAPLE achieves 11% to 242% improvement in the prediction-recall on multiple OOD benchmarks over various base models. Using MMAPLE, we reveal novel interspecies metabolite-protein interactions that are validated by activity assays and fill in missing links in microbiome-human interactions. MMAPLE is a general framework to explore previously unrecognized biological domains beyond the reach of present experimental and computational techniques.

Interactions between small molecules and proteins play pivotal roles in various biological processes across organisms. However, majority of these interactions remain understudied due to experimental constraints and human biases, limiting our understanding of the complex mechanisms governing life and hampering efforts of drug discovery.

Metabolite-protein interactions (MPIs) play a crucial role in regulating metabolic pathways, triggering signaling transduction, and maintaining cellular balance. However, MPIs are frequently low-affinity and are difficult to be detected by experiments. A recent study discovered that many overlooked MPIs contribute to the survival and growth of organisms in response to a changing environment<sup>1</sup>. Additionally, proteome-wide characterization of MPIs provides strong evidence that metabolites serve as not only intermediates in metabolic reactions but also signaling molecules via interactions with proteins that are not enzymes<sup>2,3</sup>.

In addition to intra-species MPIs, the microbiome co-evolves with the human and plays a role in shaping human phenotypes<sup>4</sup>. In the human body, microbiota produces an extremely diverse metabolite repertoire that can gain access to and interact with host cells, thus influencing the phenotype of

the human host<sup>5</sup>. For example, butyrates produced by the microbiome binds to diverse human non-enzyme proteins<sup>6</sup>. The human microbiome is not only associated with a large number of human diseases but also responsible for the efficacy and toxicity of therapeutics (e.g., cancer immunotherapy) that target human host<sup>7,8</sup>. Thus, elucidating previously unrecognized interspecies MPIs will shed light on the molecular mechanisms underlying microbiome-human interactions<sup>9,10</sup> and offer new opportunities on developing novel therapeutics<sup>11</sup>.

Beyond MPIs, uncovering novel drug-target interactions (DTIs) will facilitate identifying novel therapeutic targets, understanding poly-pharmacology, and advancing drug repurposing, thereby accelerating drug discovery and development. Unfortunately, small molecule ligands of more than 90% protein families remain unknown<sup>12</sup>. Many of these understudied proteins could be potential drug targets<sup>13</sup>. The lack of knowledge of small molecule ligands of understudied proteins hinders the drug development of presently incurable diseases<sup>14</sup>. On the one hand, many disease-causing genes are functionally and pharmaceutically uncharacterized<sup>15</sup>. It is a time-consuming and high-risk endeavor to develop assays for compound

<sup>1</sup>Ph.D. Program in Computer Science, The Graduate Center, The City University of New York, New York, NY, USA. <sup>2</sup>Department of Computer Science, Hunter College, The City University of New York, New York, NY, USA. <sup>3</sup>Helen & Robert Appel Alzheimer's Disease Research Institute, Feil Family Brain & Mind Research Institute, Weill Cornell Medicine, Cornell University, New York, NY, USA. ✉e-mail: [lei.xie@hunter.cuny.edu](mailto:lei.xie@hunter.cuny.edu)

screening of understudied proteins. On the other hand, a drug acts on a biological system, and often interacts with not only its intended target but also unknown off-target(s) that may lead to unexpected drug adverse reactions or therapeutic effects<sup>16</sup>. Thus, identifying drug-target interactions including chemicals with novel scaffolds will improve the successful rate of drug discovery for unmet medical needs<sup>17</sup>.

Given the experimental challenges in elucidating understudied intra- and inter-species molecular interactions, deep learning offers a promising alternative approach owing to its recent phenomenal successes in Natural Language Processing (NLP) and image processing. A number of deep learning models have been developed to predict drug-target interactions<sup>18–20</sup>. Nevertheless, few methods can accurately and reliably predict understudied molecular interactions due to the dearth of labeled data and out-of-distribution (OOD) problems in which small molecules or proteins involved in the interaction are significantly different from those in the annotated databases used as training data. A plethora of transfer learning techniques have been developed to address the molecular OOD problem<sup>21</sup>. However, they offer little help in bridging remote chemical spaces<sup>21</sup>.

Semi-supervised learning and meta-learning alone have shown promise in addressing OOD challenges in protein-ligand interaction predictions<sup>12,22</sup>. Semi-supervised learning uses the labeled data to learn patterns that are generalized to the unlabeled data. This approach has shown potential in exploring new chemical spaces, as demonstrated by Liu et al.<sup>22</sup>. On the other hand, meta-learning is an approach to learning to learn. It has demonstrated superior generalization power across various applications<sup>23,24</sup>. Cai et al. has developed an out-of-cluster meta-learning (OOC-ML) method, which has notably enhanced the generalization performance for OOD protein-chemical interactions<sup>12</sup>. OOC-ML simulates an OOD scenario. It harnesses common patterns extracted from predicting ligands across distinct protein clusters (meta-model) and generalizes this knowledge to a different protein cluster. These two techniques complement each other: semi-supervised learning explores unlabeled OOD data while meta-learning exploits labeled data. To our knowledge, no methods have been developed to combine these approaches to overcome data scarcity and address OOD challenges in predicting molecular interactions. Additionally, state-of-the-art semi-supervised learning method uses a teacher-student model where the teacher model is fixed in each iteration<sup>22</sup>. It may lead to confirmation bias in pseudo-labeling, another problem that needs to be addressed.

In this paper, we have developed MMAPLE - Meta Model Agnostic Pseudo Label Learning - to address the challenges aforementioned. MMAPLE incorporates the concept of meta-learning and transfer learning into a semi-supervised learning framework. Under a meta-learning framework, the student model in MMAPLE constantly sends feedback to the teacher to reduce confirmation biases. MMAPLE is effective in exploring unlabeled data and addressing the OOD problem. We have demonstrated that MMAPLE significantly improves the accuracy of DTIs, human metabolite-enzyme interactions, and understudied microbiome-human MPI predictions on multiple base models in the OOD setting. Using MMAPLE, we have identified and experimentally validated novel microbiome-human MPIs and proposed their associations with human physiology. Our findings suggest that MMAPLE can be a general framework for investigating understudied biological problems.

## Results

### Overview of MMAPLE

We evaluate the proposed MMAPLE method on three diverse OOD cases: novel DTIs, hidden human MPIs, and understudied microbiome-human MPIs. The statistics of training/validation and testing data are shown in Fig. 1A. In brief, for OOD DTIs and human MPIs, no chemicals in the testing data have a Tanimoto coefficient larger than 0.5 compared with those in the training/validation set. Details of the distribution of chemical similarities between training/validation and testing data are shown in Fig. 1C. Although 1.7% chemicals in the testing data are similar to those in the training/validation data of microbiome-human MPIs with a Tanimoto

coefficient larger than 0.5, the proteins are significantly different based on e-value as shown in Fig. 1C lower-right. Furthermore, there are no labeled microbiome-human MPIs in the training data. Thus, all benchmarks are in a challenging understudied label scarcity and/or OOD scenario.

The uniqueness of MMAPLE in predicting understudied OOD molecular interactions is threefold, as shown in Fig. 1B. Firstly, MMAPLE iteratively transfers knowledge from observed molecular interactions to the unexplored chemical genomics space of interest, employing a teacher-student approach. Secondly, unlike a conventional teacher-student model, the teacher model receives feedback from the student model to perform a meta-update aligned with meta-learning. Lastly, akin to transfer learning, the training of the student model incorporates a new target domain sampling strategy. Its aim is to guarantee that the unlabeled target domain of interest mirrors the distribution of the labeled source domain as well as to increase sampling efficiency. This alignment facilitates the model in acquiring a more robust and generalizable representation of the data. In contrast, utilizing untargeted random sampling for unlabeled data during training may lead to a markedly divergent data distribution from that of the target domain, owing to the astronomical chemical genomics space. For the three cases in this study, the target domain is different, as defined such that (1) the data distribution is significantly different from the labeled source domain to avoid data leaking, and (2) the data is relevant to the problem of interest. For example, for the DTI prediction, only drug targets are sampled.

We train base binary classification models that use labeled molecular interactions from ChEMBL<sup>25</sup>. The base models used in this study included four state-of-the-art models for chemical-protein predictions: pre-trained protein language model DISAE<sup>26</sup>, TransformerCPI<sup>18</sup>, DeepPurpose<sup>19</sup>, and BACPI<sup>20</sup>. Then MMAPLE was applied to these models for exploring unlabeled molecular interaction space. MMAPLE first initializes a teacher model using the labeled data. Then a target domain sampling strategy is applied to select a set of unlabeled data from the large space of understudied OOD DTIs or MPIs. The pre-trained teacher model makes predictions about the selected unlabeled data and assigns labels to them (pseudo labels). Next, a student model is trained using pseudo-labeled data. Different from conventional teacher-student model, the student model is evaluated by labeled data and provides feedback (metadata) to the teacher model during training the student model. Finally, the teacher model is updated based on the performance of the student and generates new pseudo labels. This process repeats multiple times until the training converges. The number of iterations depend on the base model and the problem of interest. The details of MMAPLE are in the Method section.

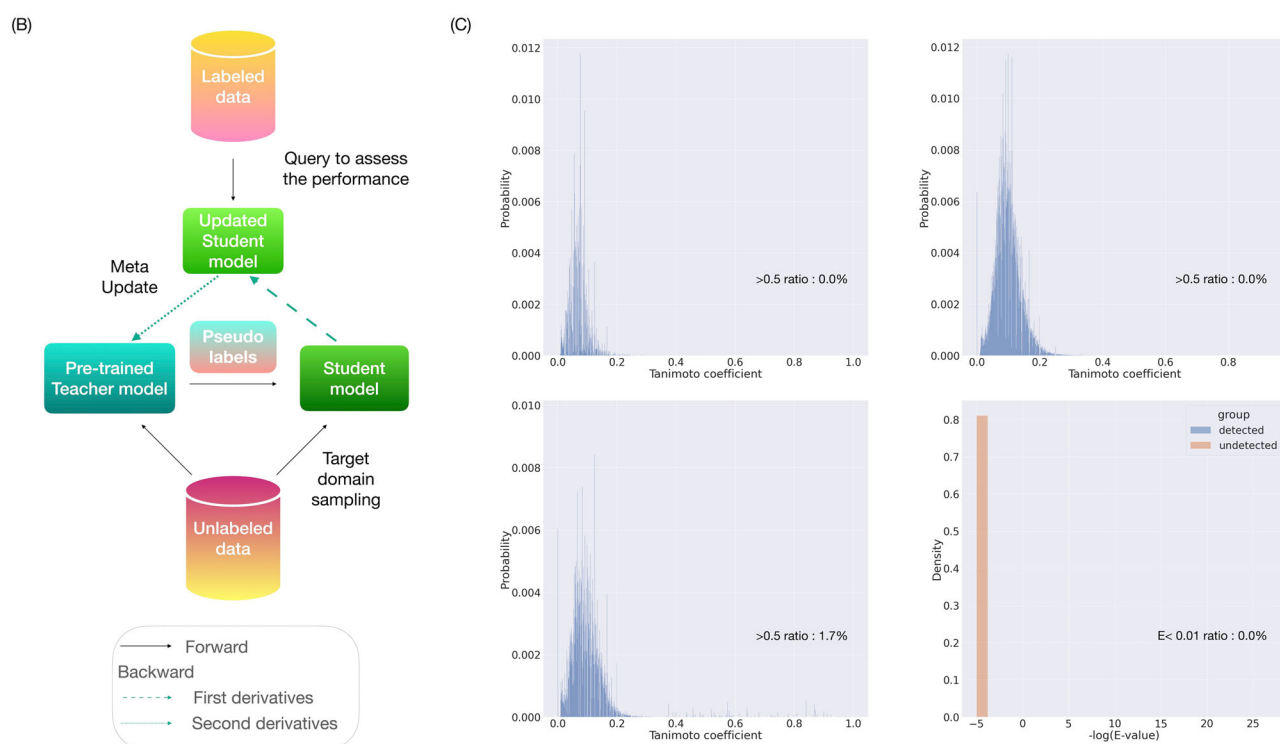
### MMAPLE significantly improves the performance of OOD DTI predictions

We first evaluated the performance of MMAPLE for OOD DTI predictions. We used molecular interactions in ChEMBL<sup>25</sup> and HMDB<sup>27</sup> as training data, and annotated DTIs from DrugBank<sup>28</sup> as testing data. To simulate an OOD scenario, we removed all chemicals that are structurally similar to drugs in the testing data (Tanimoto coefficient >0.5). As shown in Fig. 2, both PR-AUCs and ROC-AUCs of MMAPLE are significantly improved over all base models with p-values less than 0.05 (Supplemental Tables 1, 2). The percentage of improvement on PR-AUC ranges from 13% to 26%. Furthermore, the trained models are less over-fitted than the base models, as supported by the narrow gaps between the training curve of validation data and that of testing data, as shown in Supplemental Fig. 1.

The superior performance of MMAPLE may be because it can better align the embedding space of OOD samples to that of training data. To test this hypothesis, we investigated if MMAPLE could alleviate the distribution shift between training and testing data. We extracted the embeddings of the training and testing examples before training and acquired by DISAE and MMAPLE, then utilized the Uniform Manifold Approximation and Projection (UMAP) for visual analysis. Figure 2C supports this hypothesis. Before the training, the embeddings of training chemicals are scattered around those of testing chemicals. While DISAE - the best-performed base

(A)

Experiment	Sample size			Note		
OOD DTI	train/val	unique chemicals	230,811	Coverage of labeled chemical-protein pairs	Target domain sampling	Distribution of chemical similarity (Tanimoto)
		unique proteins	3,084			
	total pairs		298,736			
	test	unique chemicals	8,917			
unique proteins		3,266				
unlabeled data pairs used in training			53,502	0.04%	for each chemicals, sample six proteins	>0.5: 0.0%
Hidden human MPI	train/val	unique chemicals	252,712	0.03%	for each chemicals, sample five proteins	>0.5: 0.0%
		unique proteins	5,204			
	total pairs		334,668			
	test	unique chemicals	8,921			
unique proteins		2,611				
unlabeled data pairs used in training			44,644			
Zero-shot microbiome-human MPI	train/val	unique chemicals	357,213	0.00%	proteins: GPCR + the same number of proteins from Uniprot; chemicals: NJS16	>0.5: 1.7%
		unique proteins	168,517			
	total pairs		1,667,708			
	test	unique chemicals	109			
unique proteins		65				
unlabeled data pairs used in training			73,238			



**Fig. 1 | Data statistics and framework illustration.** **A** Statistics of training, validation, and testing data used in this study. **B** Illustration of MMAPLE framework. A deep learning model is trained using both labeled and unlabeled data and iteratively updated using gradients from the trained model as metadata. **C** Chemical and protein similarity distribution between training/validation and testing datasets. Up-

left: Chemical similarity distribution of OOD DTI; up-right: chemical similarity distribution of hidden human MPI; lower left-right: chemical and protein similarity distribution of zero-shot microbiome-human MPI experiment. Chemical similarity is quantified by the Tanimoto coefficient of chemical fingerprints. Protein similarity is measured by the negative logarithm (base 10) of the e-value derived from BLAST<sup>61</sup>.

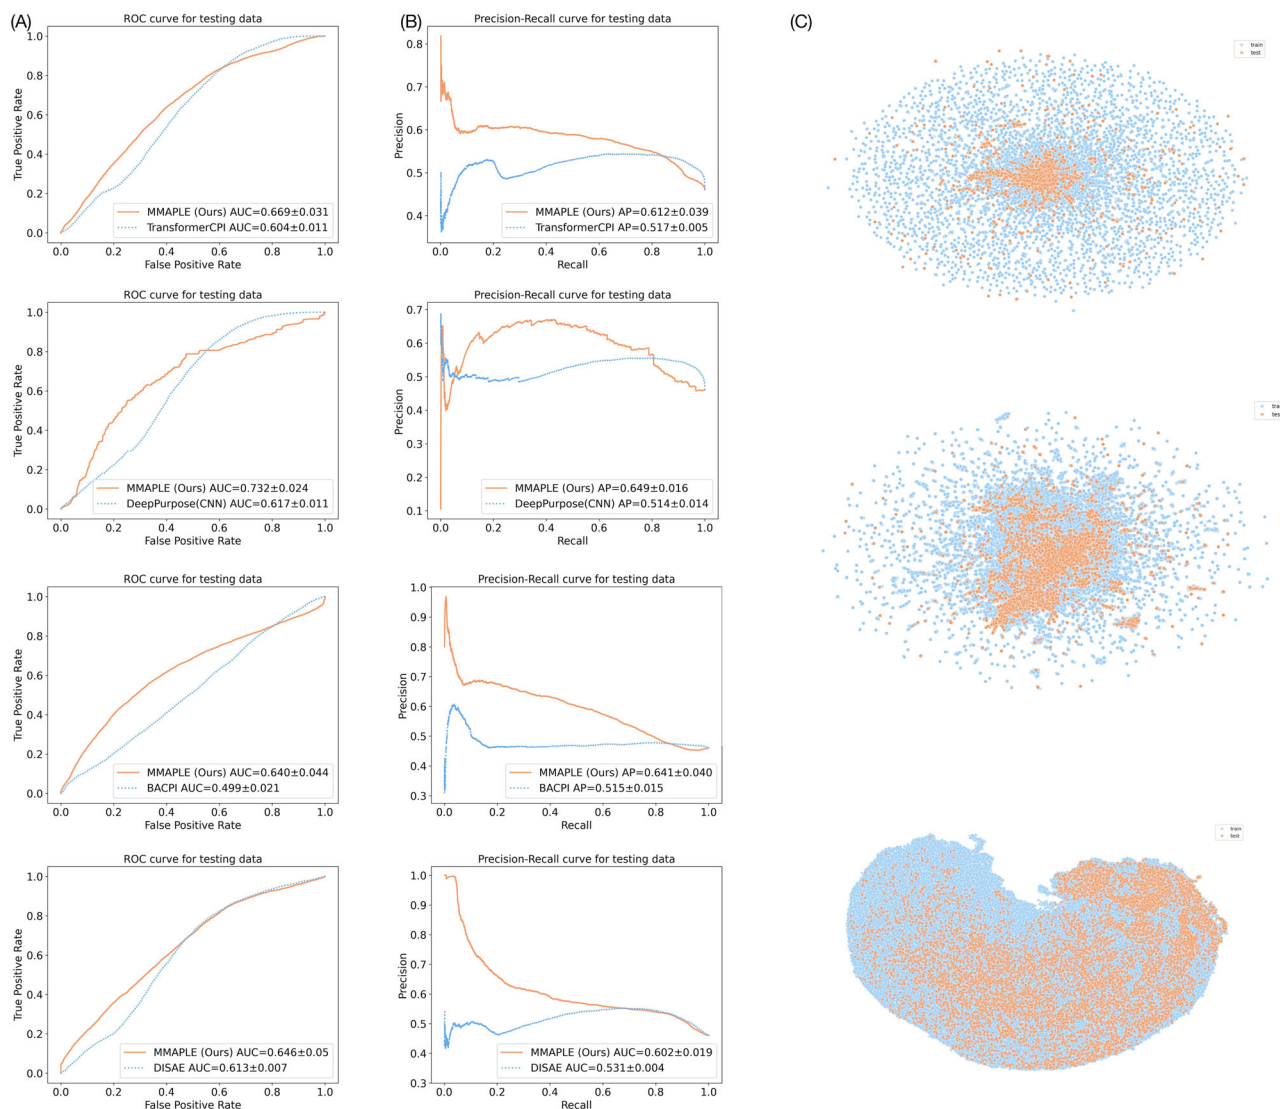
model - narrows the dispersion, our model achieves tighter overlap between two distributions. Importantly, our model not only draws them closer but also ensures a more uniform distribution within each group, reducing inter-distribution gaps.

Transfer learning in the protein space via protein language modeling can improve the performance of DTI prediction<sup>26</sup>. As shown in Figs. 2A, B and 3A, B in the next section, DISAE that is based on a pre-trained protein language model outperforms other baselines that do not utilize the language model. However, the improvement from DISAE is not as significant as that by MMAPLE. Additionally, We study if transfer learning in the chemical space could boost the performance of OOD DTI predictions. We apply a chemical pretraining-fine-tuning based on self-supervised Motif Learning

Graph Neural Network (MoLGNN)<sup>29</sup>. In consistent with recent findings<sup>21</sup>, no improvement was detected, as shown in Supplemental Table 3.

### MMAPLE significantly improves the performance of hidden OOD human MPI predictions

We next evaluated the performance of the MMAPLE model in predicting hidden human MPIs. We first trained the model using ChEMBL which primarily includes exogenous small molecule ligands and druggable protein targets. We evaluated the performance of the trained model using the Human metabolite database HMDB<sup>27</sup> on human MPIs. The test cases were in the OOD setting, as supported by the chemical similarity distributions (Fig. 1C).



**Fig. 2 | OOD DTI prediction outcomes when applying MMAPLE to base models. A** ROC curves. **B** PR curves. **C** UMAP visualization of chemical space. Top to bottom - original fingerprints, baseline embeddings from DISAE, and MMAPLE embeddings. Orange and blue dots for MMAPLE and baseline, respectively.

Figure 3 A indicates that MMAPLE significantly outperforms all of state-of-the-art base models on both ROC and PR. The ROC-AUC and PR-AUC increase by 17% to 20% and 17% to 30%, respectively, suggesting that MMAPLE is able to accurately predict hidden human MPIs in an OOD setting.

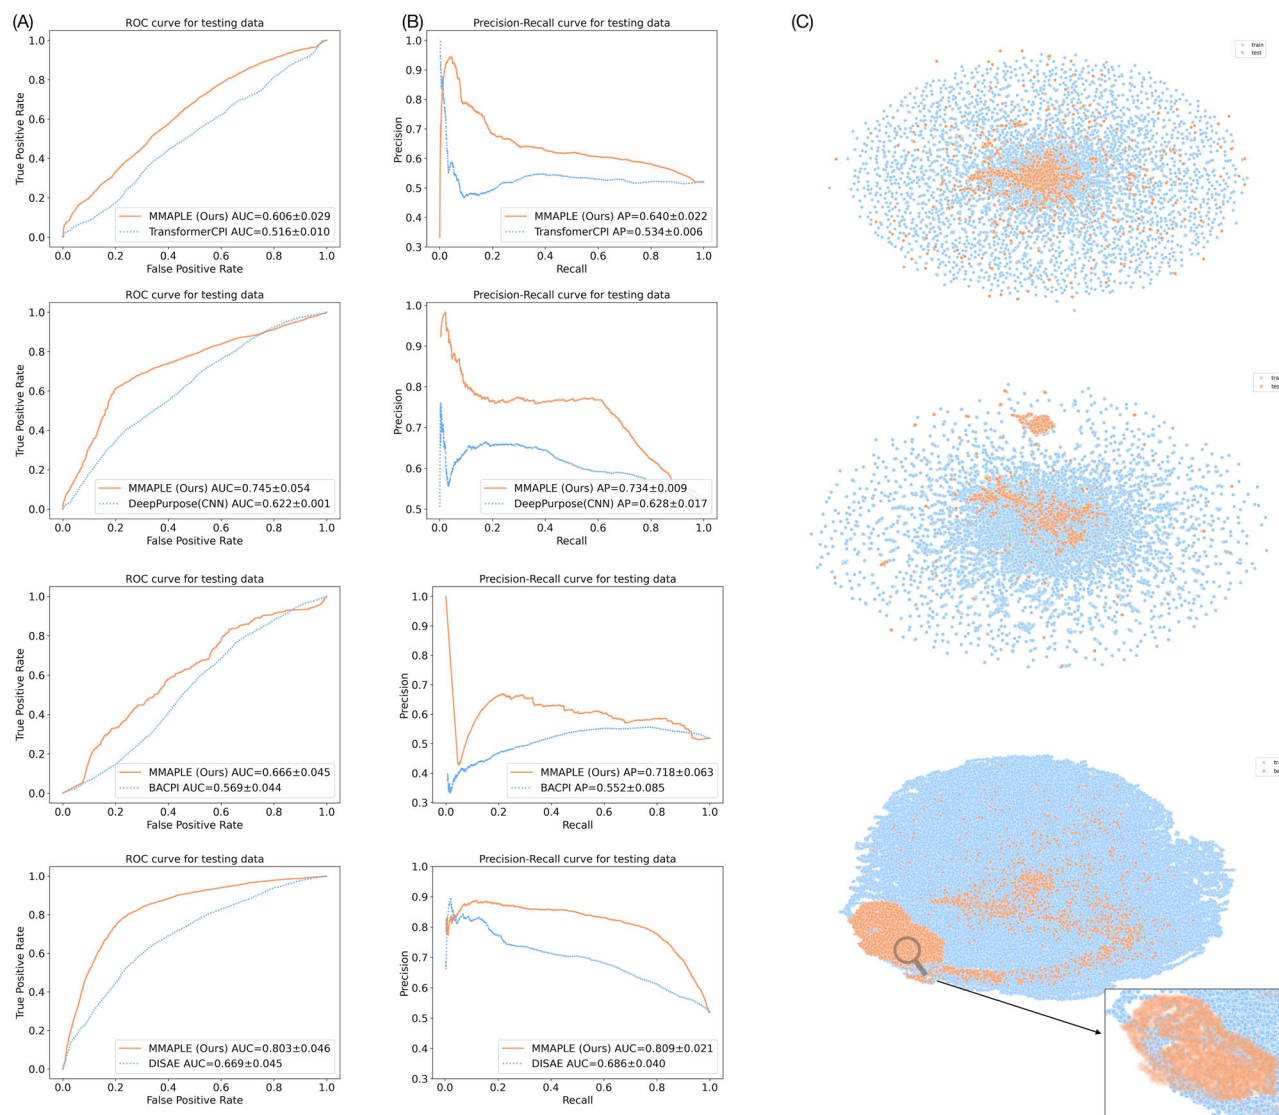
Again, MMAPLE training brings the embeddings of testing samples closer to those of training data than the baseline, as shown in Fig. 3C. Overall, our results suggest that MMAPLE significantly outperforms the state-of-the-art methods for OOD DTI and hidden MPI predictions.

### MMAPLE significantly improves the performance of understudied OOD interspecies MPI predictions and reveals the molecular basis of microbiome-human interactions

Known interspecies microbiome-human MPIs are extremely scarce, only including 17 observed active interactions (See Methods for details). To investigate interspecies interaction, MMAPLE was trained on a combination of three datasets: HMDB, ChEMBL, and NJS16<sup>30</sup>, while the test set consisted of 17 annotated along with 145 negative microbiome-human MPIs from the literature<sup>31,32</sup>. As shown in Fig. 1C, no metabolite-protein pairs in the testing set have similar chemicals or proteins to those in the training/validation set. Because no interspecies MPIs exist in the training and validation set, the problem is a zero-shot learning scenario. The

previously best performed model DISAE has a poor performance, with a PR-AUC of 0.193. It indicates that transfer learning alone is not sufficient to address the OOD challenge of interspecies MPI predictions. Our results, presented in Fig. 4A, demonstrate that MMAPLE significantly outperforms DISAE in terms of ROC and PR. It achieves a three-fold increase in the PR-AUC for interspecies MPI predictions. These findings indicate that MMAPLE holds promise in deepening our comprehension of interspecies interactions, thus serving as a valuable tool for investigating the impact of the microbiome on human health and disease.

To further validate the performance of MMAPLE, we predicted and experimentally validated the interactions between trimethylamine N-oxide (TMAO) and human G-protein coupled receptors (GPCRs). TMAO is a small molecule generated by gut microbial metabolism. It has been observed that elevated plasma levels of TMAO increase the risk for major adverse cardiovascular events<sup>33</sup>, activate inflammatory pathways<sup>34</sup>, and promote foam cell formation<sup>35</sup>. Additionally, TMAO inhibits insulin signaling<sup>36</sup>. However, it remains elusive how TMAO modulates these pathological processes at a molecular level. Besides its biological interest, TMAO is one of the most challenging molecules for MMAPLE. Firstly, the current study of microbiome-human interactions mainly focuses on short-chain fatty acids, there are few data for TMAO. Secondly, TMAO is a molecule with different structural characteristics from other chemicals in the training data, The



**Fig. 3 | Hidden Human MPI prediction outcomes when applying MMAPLE to base models. A** ROC curves; **B** PR curves; **C** UMAP visualization of chemical space. top to bottom - original fingerprints, baseline embeddings from DISAE, and MMAPLE embeddings. Orange and blue dots for MMAPLE and baseline, respectively.

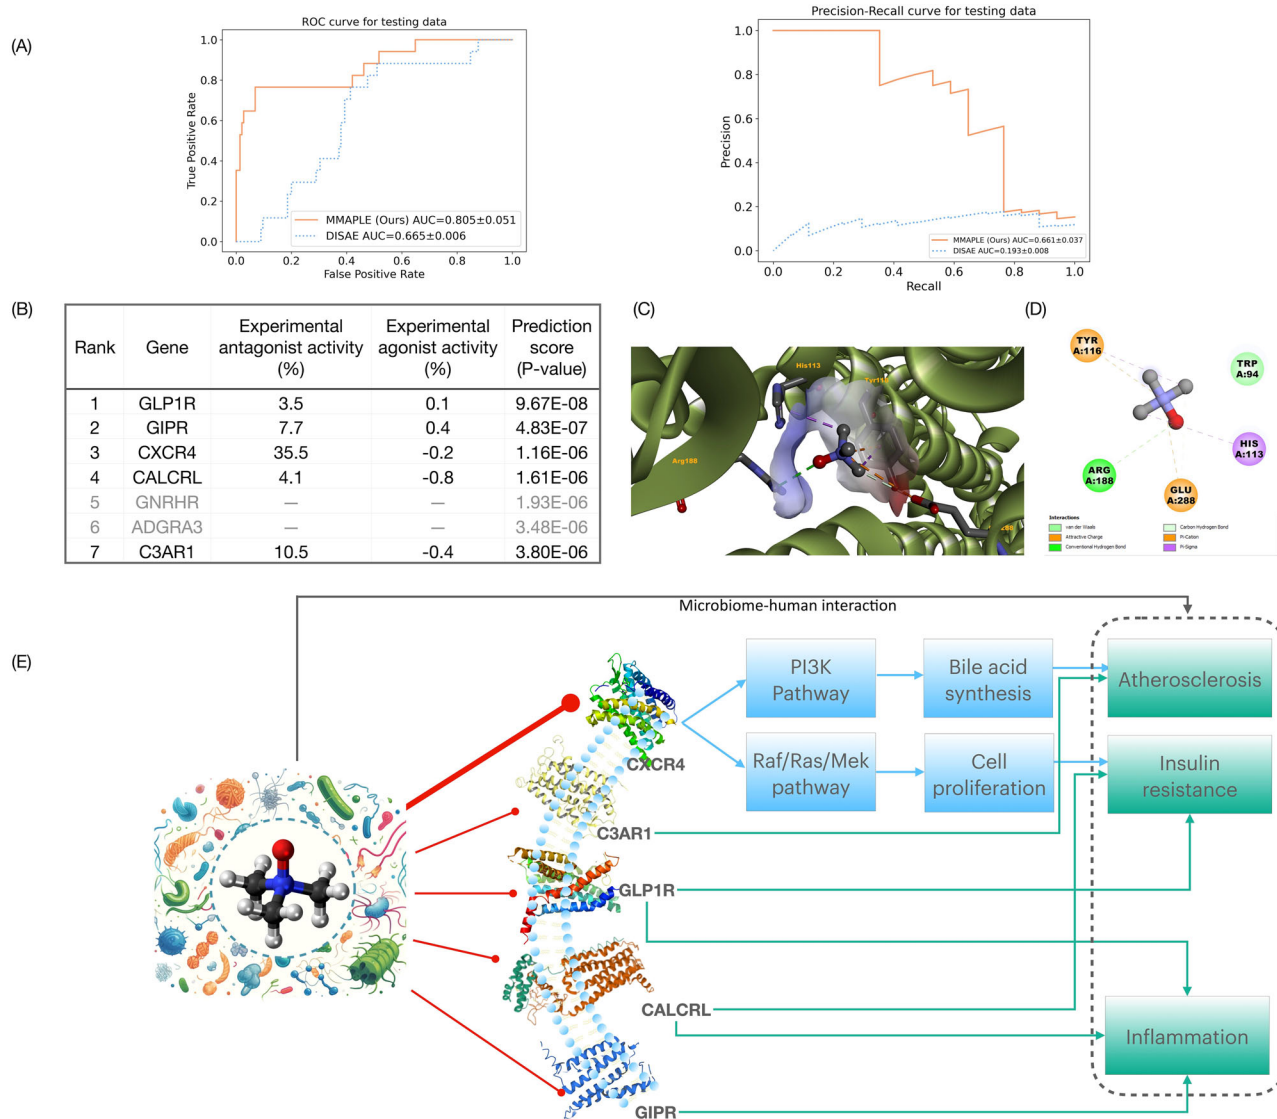
chemical structure of TMAO is significantly different from known metabolites involved in microbiome-human MPis, as shown in Supplemental Fig. 2. Thus, we choose TMAO to rigorously evaluate MMAPLE in an OOD scenario.

Figure 4B lists the top 7 predicted GPCR genes that interact with TMAO with a p-value less than  $5.0 \times 10^{-6}$  (approximate false discovery rate of 0.05). We performed GPCR functional assays to experimentally test the binding activities of five of them under the concentration of  $30 \mu\text{M}$  of TMAO, which is the same concentration used in the previous study and is based on the physiological concentration of TMAO in the human ( $1\text{--}45 \mu\text{M}$ )<sup>37</sup>. The assay for two top-ranked GPCRs *GNRHR* and *ADGRA3* is not available. As shown in Fig. 4B, all five tested GPCRs are antagonists that block the activity of receptors, and *CXCR4* demonstrates the strongest activity (activity score  $>30$ ). Other top-ranked predictions can be found in Supplemental Table 4. We have also performed additional experiments to analyze the predictions from the baseline model (DISAE). MMAPLE significantly outperforms the base model, as shown in Supplemental Figure 3. The full predictive results are in Supplemental Table 5.

Protein-ligand docking by AutoDock Vina<sup>38</sup> suggests that TMAO can fit into the antagonist conformation of the *CXCR4* structure, as shown in Fig. 4C, D. Among these interacting residues, TRP94, TYR116, and Glu288 also interact with the co-crystallized ligand of encoded protein of *CXCR4*

(PDB id: 3ODU). TYR116 and GLU288 provide attractive charges to the nitrogen atom on TMAO. ARG188 forms a conventional hydrogen bond with an oxygen atom on TMAO. These strong interactions could keep TMAO in the binding pocket. The *CXCR4* antagonism by TMAO establishes a causal linkage for observed microbiome TMAO-human interactions, as illustrated in Fig. 4E. It is known that *CXCR4* regulates PI3K and RAF/RAS/MEK pathways<sup>39</sup> (KEGG Pathway: <https://www.genome.jp/pathway/hsa04062>). PI3K pathway regulates bile acid synthesis<sup>39</sup>. TMAO's inhibition on bile acid synthesis may be responsible for its promotion effect on atherosclerosis<sup>33</sup>. The physiological effect of TMAO on obesity and insulin resistance may be via *CXCR4*-RAF/RAS/MEK axis. It has been observed that the deficiency of *CXCR4* and impaired RAF/RAS/MEK signaling results in obesity and insulin resistance<sup>40–43</sup>. Thus, microbiome TMAO-human *CXCR4* interaction is responsible for the several observed pathological effects of TMAO. However, other TMAO effects such as inflammation cannot be directly explained by the TMAO-*CXCR4* interaction. It is possible that other human proteins can interact with TMAO.

While other top-predicted GPCRs (*GLP1R*, *GIPR*, *CALCRL*, *C3AR1*) show much weaker binding to TMAO compared to *CXCR4* at the tested TMAO concentration, the antagonist (inhibition) effect of TMAO can be amplified at higher TMAO concentrations, such as when consuming food,



**Fig. 4 | Results of microbiome metabolite-human protein interaction predictions.** **A** ROC and PR curves of models tested by literature annotated 17 positive and 145 negative microbiome-human MPis. **B** Top 5 predicted G-protein coupled receptor (GPCR) genes that interact with TMAO and GPCR functional assay results, with p-value indicating the tail probability from Kernel density estimation; **C** Predicted 3D binding pose of TMAO on the CXCR4 antagonist conformation;

**D** Interaction patterns between TMAO and CXCR4; **E** Proposed molecular mechanism of TMAO-human interactions. No assay is available for GNRHR and ADGRA3. Antagonist: a molecule blocks the receptor, preventing activation and the usual cellular response. Agonist: a molecule activates the receptor, triggering a cellular response.

and aligns with experimental evidence. *GLPIR* activation is known to reduce inflammation, suggesting that TMAO inhibition of *GLPIR* might increase inflammation<sup>44,45</sup>. Similarly, studies have shown that gut *GIPR* is associated with diet-induced inflammation and insulin resistance<sup>46</sup>. Modulation of *CALCRL* is associated with insulin resistance<sup>47</sup>, and its deletion can worsen intestinal inflammation<sup>48</sup>. *C3AR1* plays a protective role against atherosclerosis<sup>49</sup>, implying that TMAO blocking *C3AR1* activity might increase the risk of atherosclerosis in individuals.

### Semi-supervised learning, meta-learning, and target domain sampling synergistically contribute to the performance of MMAPLE

In our comprehensive ablation study, we rigorously examined the influence of several key components on our model’s performance: the introduction of target domain sampling and semi-supervised learning with pseudo labels, the choice between utilizing soft pseudo labels or hard labels, and the application of meta-learning.

When excluding meta-learning from our training process, we kept the teacher model static, therefore restricting it to generate constant pseudo labels for the student model to learn from. This leads to performance decline when compared to our full MMAPLE model, as shown in Table 1. The absence of iterative feedback learning addresses the critical role of meta-learning.

To investigate the effect of teacher-student training, we trained the model solely with meta-learning by leveraging the Model-Agnostic-Meta-Learning (MAML) framework<sup>50</sup>. This approach, while constantly outperforming the baseline, resulted in 124% fall in PR-AUC compared to MMAPLE. This experiment not only demonstrated the intricate dependencies between meta-learning and semi-supervised learning but also underscored the necessity of synergy of these techniques to achieve superior model performance.

Models trained on one-hot (hard) labels are subject to over-fitting since they do not represent soft decision boundaries across concepts. Soft labels, which are probability distributions over the possible classes

**Table 1 | Ablation study**

Methods	ROC-AUC	PR-AUC
base model (DISAE)	0.665 ± 0.006	0.193 ± 0.008
meta + TS soft	0.727 ± 0.033	0.308 ± 0.020
meta + Vanilla TS	0.708 ± 0.013	0.247 ± 0.024
OOO-ML (out-of-cluster meta-learning)	0.748 ± 0.009	0.295 ± 0.034
TS soft + target	0.791 ± 0.061	0.632 ± 0.108
TS soft	0.655 ± 0.065	0.209 ± 0.001
Vanilla TS	0.665 ± 0.059	0.198 ± 0.006
<b>meta + TS soft + target (MMAPLE)</b>	<b>0.805 ± 0.051</b>	<b>0.661 ± 0.037</b>

The variance was derived through k-fold cross-validation with k = 3

as opposed to hard labels, are often demonstrated to be more effective due to the ability to provide the model with more information about the uncertainty in the data, as well as the ability against label noise, resulting in more robust predictions<sup>51,52</sup>. As shown in the Table 1, when soft labels were used, ROC-AUC improved by 25%, and PR-AUC increased by twofold.

The objective of target domain sampling aligns with that of transfer learning. As shown in Table 1, using target domain sampling significantly increased the performance of the model by 11% on ROC-AUC and 115% on PR-AUC, showing the effectiveness of this strategy in improving the performance of MMAPLE.

In summary, integrating meta-learning, target domain sampling, and soft labeling into a teacher-student framework yields superior performance compared to each of these approaches individually, as well as any combination of two of them.

## Discussion

In this study, we present MMAPLE, a highly effective deep learning framework, designed to address the challenges of data scarcity and OOD problems encountered when applying machine learning in understudied biological domains when transfer learning is less effective. Through extensive evaluations, we have demonstrated the exceptional capabilities of MMAPLE in exploring the unlabeled data space and facilitating knowledge transfer from one chemical space to another. Using MMAPLE, we successfully predicted and experimentally validated novel interactions between microbiome metabolites and human proteins, thereby shedding light on the intricate interplay between these components. Notably, our framework does not rely on a specific model and can accommodate various deep-learning architectures tailored to specific biological tasks. Thus, MMAPLE serves as a versatile and robust framework for investigating a wide range of understudied biological problems.

MMAPLE shows potential for improvement in several key areas. Firstly, the current implementation of MMAPLE lacks the ability to estimate the uncertainty associated with pseudo labels. By incorporating an accurate uncertainty quantification mechanism, it becomes possible to select high-confidence pseudo labels during training, therefore reducing the impact of noise. Secondly, the process of sampling pseudo labels in a vast and imbalanced chemical-protein interaction space proves time-consuming, particularly when aiming to achieve the desired positive versus negative ratio. The performance of MMAPLE can be further enhanced by employing an unbiased and efficient sampling strategy. For example, sampling based on protein family or chemical similarity clustering. Thirdly, while MMAPLE has thus far been applied exclusively to classification problems, it would be interesting to explore its extension to regression problems. Lastly, the meta-update in the current implementation of MMAPLE uses indistribution data. Incorporation of OOC meta-learning may further improve the generalization power of MMAPLE in an OOD setting. These would be subject to future study.

## Method

### Data sets

#### Experiment 1: DTI prediction

Training/validation data. We used molecular interactions in ChEMBL<sup>25</sup> and HMDB<sup>27</sup> as training and validation data. It contained 298,736 total pairs with 230,811 unique chemicals and 3084 unique proteins.

OOD testing data. The annotated DTIs from DrugBank<sup>28</sup> were used as testing data. To simulate an OOD scenario, we removed all chemicals that are structurally similar to drugs in the testing data (Tanimoto coefficient > 0.5), totaling 21,760 pairs including 8917 unique chemicals and 3266 proteins.

Unlabeled data. To focus on the unexplored domain of interest, a target domain sampling strategy was developed. Specifically, we selected unlabeled pairs of drug targets and drug-like chemicals but excluded already labeled pairs. For each chemical, we sampled six proteins, resulting in 53,502 total unlabeled pairs. The detailed data statistics can be found in Fig. 1.

#### Experiment 2: Human MPI prediction

Training/validation data. The training data for this experiment was sourced from ChEMBL (version 29)<sup>25</sup>. It consisted of 334,668 pairs with 252,712 unique compounds and 5204 unique proteins, where each pair represented an activity with a single protein as the target.

OOD testing data. For the testing, we utilized HMDB<sup>27</sup>, which provided interactions between metabolites and human enzymes. We randomly sampled 10,000 pairs as the testing data covering 8,921 unique compounds and 2611 unique proteins.

Unlabeled data. To create the unlabeled dataset, we considered all the unlabeled metabolite-enzyme pairwise combinations. From the total pairs, we included all unique metabolites and randomly selected two enzymes to associate with each chemical, this resulted in the creation of a sizeable unlabeled dataset, consisting of 44,644 unlabeled samples. The detailed data statistics can be found in Fig. 1.

#### Experiment 3: Microbiome-human MPI prediction

Training/validation data. For this experiment, the training data consisted of a combination of ChEMBL, HMDB, and NJS16<sup>30</sup> datasets. After removing duplicates and unusable data, the dataset contained a total of 1,667,708 samples including 357,213 unique compounds and 168,517 unique proteins.

OOD testing dataset. The testing dataset was manually created based on two published works. The first work<sup>31</sup> provided information on interactions between 241 GPCRs and metabolites from simplified human microbiomes (SIHUMIs) consisting of the seven most general bacteria species. The second work<sup>32</sup> involved the screening of gut microbiota metabolomes to identify ligands for various GPCRs. Since this study focused on small molecule metabolites, lipids were excluded, resulting in a total of 162 MPIs, including 17 positive activities.

Unlabeled data. For the protein side, we included all GPCRs from UniProt<sup>53</sup>. Besides, an equal number of proteins were randomly selected from the Pfam dataset. Chemical samples were the 240 unique metabolites from the NJS16 dataset. Overall, the unlabeled data consisted of 73,238 pairs. The detailed data statistics can be found in Fig. 1.

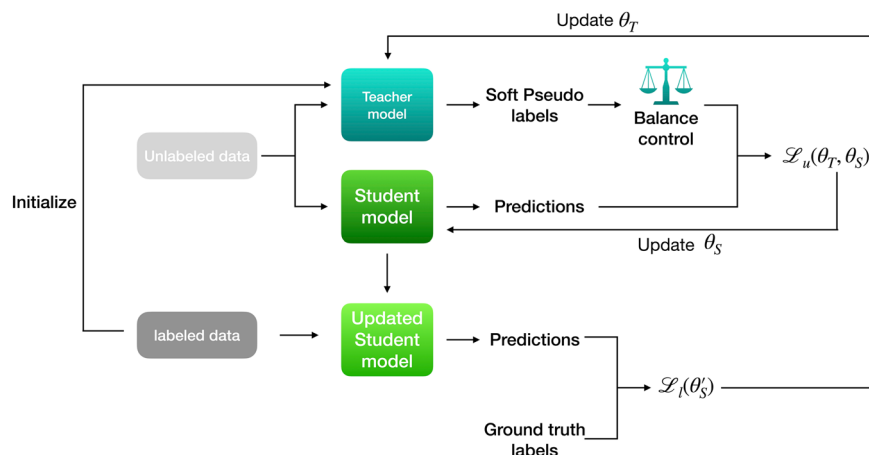
### MMAPLE base models

To enable a fair comparison with the baseline models, we currently focus on binary classification problems. Four state-of-the-art base models were employed to evaluate the performance MMAPLE:

- DISAE<sup>26</sup>. Distilled Sequence Alignment Embedding (DISAE) is a method developed by us that includes three major modules: protein language model, chemical structure modeling, and the combination of

**Fig. 5 | Illustration of MMAPLE training schema.**

A teacher model generates pseudo labels by predicting a batch of unlabeled data. The pseudo label is further passed to a filter to control the balance ratio of the positive and negative samples (as a hyper-parameter). A student model generates the predictions from the same unlabeled data as those used in the teacher model and is updated by minimizing loss function  $\mathcal{L}_u(\theta_T, \theta_S)$  as in equation (3). Then, the updated student model takes a batch of labeled data and generates new predictions that compare with the ground truth labels and minimize loss  $\mathcal{L}_l(\theta'_S)$  as equation (6).



the above two modules. The protein sequence module uses distilled sequence alignment embedding, leveraging a transformer-based architecture trained on nearly half a million protein domain sequences for generating meaningful protein embeddings. This is crucial for predicting protein-ligand interactions in out-of-distribution (OOD) scenarios. The chemical module is a graph isomorphism network (GIN) to obtain chemical features, which are numerical representations of small molecules and capture their chemical properties. Finally, DISAE includes an attentive pooling module that combines the protein and chemical embeddings obtained from the first two modules to produce the final output for predicting DTIs or MPIs as a binary classification task (i.e., active or inactive). The attentive pooling module uses a cross-attention mechanism to weigh the importance of each protein and chemical embedding, allowing it to focus on the most relevant information when making the prediction.  $L_{base}$  denotes the loss function of the base model, which is a binary cross-entropy loss in this case.

- TransformerCPI<sup>18</sup> Adapted from the transformer architecture, TransformerCPI takes protein sequence as the input to the encoder, and atom sequence as the input to the decoder, and learns the interaction at the last layers. Specifically, the amino acid sequence is embedded with a Word2vec model pre-trained on all human protein sequences in UniProt, and the self-attention layers in the encoder are replaced with a gated convolutional network and output the final presentation of proteins. The atom features of chemicals are learned through graph convolutional network (GCN) by aggregating their neighbor atom features. The interaction features are further obtained by the decoder of transfer, which consists of self-attention layers and feed-forward layers.
- DeepPurpose<sup>19</sup> DeepPurpose provides a library for DTI prediction incorporating seven protein encoders and eight compound encoders to learn the protein and compound representations respectively, and eventually feeds the learned embeddings into an MLP decoder to generate predictions. We implemented the best-reported architecture, convolutional neural network (CNN) for both protein and compound feature representation learning, as another base model of MMAPLE.
- BACPI<sup>20</sup> The last base model included in this study is the Bi-directional attention neural network for compound-protein interaction (BACPI). Similarly, it consists of chemical representation learning, protein representation learning, and CPI prediction components to combine them. BACPI employs a graph attention network (GAT) for compounds to learn various information of the molecule graphs. For protein, it introduces a CNN module to take the amino acid sequence as input, to learn the local contextual features of protein by using a content window to split the sequences onto overlapping subsequences

of amino acids. Finally, the atom structure graphs and residue sequence features are fed into the bi-directional attention neural network to integrate the representations and capture the important regions of compounds and proteins, and the integrated features are used to predict the CPI.

### Semi-supervised meta-learning

We adopted a semi-supervised meta-learning paradigm for our model training. Similar to pseudo labels, there is a pair of teacher model and student model, the teacher model takes unlabeled data as input, and uses the predicted results as pseudo labels for the student model to learn with the combination of labeled and pseudo-labeled data. However, instead of learning from the fixed teacher model, the student constantly sends feedback to the teacher in the format of performance on labeled data, and the teacher keeps updating the pseudo labels on every mini-batch. This strategy could solve the problem of confirmation bias in pseudo-labeling<sup>54</sup>. The illustration of MMAPLE training is shown in Fig. 5. Let  $T$  and  $S$  denote the teacher model and the student model,  $\theta_T$  and  $\theta_S$  denote the corresponding parameters ( $\theta'_T$  and  $\theta'_S$  denote the updated parameters). We use  $\mathcal{L}$  to represent the loss function, and  $T(x_u; \theta_T)$  to stand for the teacher predictions on unlabeled data  $x_u$ , similar notations for  $S(x_u; \theta_S)$  and  $S(x_l; \theta'_S)$ .  $CE$  denotes the cross-entropy loss.

### Model training

MMAPLE does not work alone but is built on top of other models. The training process is repeated until optimization converges. The number of iterations depends on the base model and training data, so it varies accordingly. To ensure a fair comparison with the base models, both MMAPLE and base models were constructed using the same architecture. The detailed training procedure is shown in Algorithm 1.

**The update rule of student.** On a batch of unlabeled data  $x_u$ , sample  $T(x_u; \theta_T)$  from the teacher's prediction, and optimize the student model with the objective

$$\min_{\theta_S} \mathcal{L}_u(\theta_T, \theta_S) \quad (1)$$

where

$$\mathcal{L}_u(\theta_T, \theta_S) := \mathbb{E}_{x_u} [CE(T(x_u; \theta_T), S(x_u; \theta_S))] \quad (2)$$

The optimization of each mini-batch is performed as

$$\theta'_S = \theta_S - \eta_S \nabla_{\theta_S} \mathcal{L}_u(\theta_T, \theta_S) \quad (3)$$



**The update rule of teacher.** On a batch of labeled data  $(x_i, y_i)$ , and use the students' update to optimize the teacher model with the objective

$$\min_{\theta_T} \mathcal{L}_l(\theta_S - \eta_S \nabla \theta_S \mathcal{L}_u(\theta_T, \theta_S)) \quad (4)$$

where

$$\mathcal{L}_l(\theta'_S) := \mathbb{E}_{x_i, y_i} [CE(y_i, S(x_i; \theta'_S))] \quad (5)$$

The optimization of each mini-batch is performed as

$$\theta'_T = \theta_T - \eta_T \nabla \theta_T \mathcal{L}_l(\theta_S - \eta_S \nabla \theta_S \mathcal{L}_u(\theta_T, \theta_S)) \quad (6)$$

We experimented with both hard labels and soft labels. Due to the superior performance of soft labels to hard labels, the final MMAPLE was trained using the soft label. The methods are described as follows:

**Using soft labels.** Because we always treat  $\theta_S$  as fixed parameters when optimizing Equation (6) and ignore its higher-order dependence on  $\theta_T$ , the objective is fully differentiable with respect to  $\theta_T$  when soft pseudo labels are used, i.e.,  $T(x_u; \theta_T)$  is the full distribution predicted by the teacher model. This allows us to perform standard back-propagation to obtain the gradient.

Additionally, we incorporated the temperature scaling to soften the teacher model's predictions<sup>55</sup>.  $T(x_u; \theta_T)$  is the teacher's output distribution computed by applying softmax over the logits  $z$ :  $\text{softmax}(z) = \frac{\exp(z/T)}{\sum_{j=1}^r \exp(z_j/T)}$ ,

the temperature parameter  $T$  is used to control the "softness" of the output probabilities. In the implementation, the temperature was tuned by hyperparameter searching.

For the quality control of soft labels, we employed a balance sampler to control the ratio between positive and negative hard labels transferred from soft labels. This will provide a mechanism to dynamically adjust the ratio of positive and negative during training. This ratio served as a crucial parameter to govern the training process, enabling us to strike a balance between the two label categories. Through this approach, we aimed to alleviate bias and imbalance in the dataset.

**Using hard labels.** When using hard pseudo labels, we followed the derivative rule proposed in the reference<sup>54</sup>, which was a slightly modified version of REINFORCE applied to obtain the approximated gradient of  $\mathcal{L}_l$  in Equation(6) with respect to  $\theta_T$  as follows:

$$h = \eta_S \cdot \left( \left( \nabla_{\theta_S} CE(y_i, S(x_i; \theta_S^{t+1})) \right)^T \cdot \nabla \theta_S CE(\hat{y}_u, S(x_u; \theta_S^t)) \right) \quad (7)$$

The teacher's gradient from the student's feedback:

$$g'_T = h \cdot \nabla \theta_T CE(\hat{y}_u, T(x_u; \theta_T))|_{\theta_T} = \theta_T^{(t)} \quad (8)$$

### Algorithm 1. Training procedure

**Require:**  $N$ , the batch size  $n_{sup}$ , number of epochs of supervised training  $n_{freeze}$ , number of epochs that teacher model is frozen  $n$ , number of training epochs

**Input:**  $X_{un}$ ,  $X_l$

**Stage 1:**

**for**  $epoch = 1$  to  $n_{sup}$  **do**

**for**  $t = 1$  to  $\frac{N_l}{N}$  **do**

sample  $X_l$  of size  $N$  from the labeled data (without rep)

Update  $\theta_T$  with  $L_{base}$

**end for**

**end for**

save the model with early stopping

**Stage 2:**

Initialize the teacher model with  $\theta_T$

Initialize student model with random parameters  $\theta_S$

**for**  $epoch = 1$  to  $n_{freeze}$  **do**

**for**  $t = 1$  to  $\frac{\min(N_{un}, N_l)}{N}$  **do**

sample  $X_{un}$  of size  $N$  from unlabeled data (without rep)  
update  $\theta_S$  with student update rule

**end for**

**end for**

**for**  $epoch = n_{freeze} + 1$  to  $n$  **do**

sample  $X_{un}$  of size  $N$  from unlabeled data (without rep)

update  $\theta_S$  with student update rule

update  $\theta_T$  with teacher update rule

**end for**

### Model evaluation

The model performance was measured using both Receiver Operating Characteristic (ROC) and Precision-Recall (PR) and their corresponding area under the curve (AUC). While ROC is a commonly used metric, it may give an optimistic impression of the model's performance, particularly when datasets are imbalanced<sup>56</sup>. Therefore, PR is a better metric to evaluate the performance of MMAPLE than ROC. A three-fold cross-validation approach was utilized to ensure the robustness of the model's performance evaluation. Consistency across evaluations was maintained by using the same folds for all base models.

### Statistical significance of prediction

In our study, we focused on predicting GPCR genes that interact with TMAO, employing a comprehensive analytical approach to evaluate the statistical significance of each prediction. The prediction scores generated through our model were subjected to Kernel Density Estimation (KDE) from the Python package *scipy*<sup>57</sup>. KDE is a non-parametric way to estimate the probability density function of our prediction scores. By applying KDE, we were able to calculate the tail probability for each predicted interaction score, which we interpreted as a p-value. This p-value serves as an indicator of the rarity or significance of the predictions within the overall distribution of scores, providing a statistical basis for identifying the most significant GPCR-TMAO interactions. The detailed results of our predictions can be found in Supplemental Table 4. The results of DISAE predictions can be found in Supplemental Table 5.

### Statistics and Reproducibility

Statistical analyses in general were conducted using paired t-tests. We employed three-fold cross-validation to ensure the robustness of our results. For each fold, we applied early stopping and tested the model on the hold-out testing set. The final reported mean performance is the average result from these testing sets.

### GPCR functional assay

Trimethylamine N-oxide (TMAO) (purity: 95%, molecular weight: 76.12) was purchased from Sigma-Aldrich (MO, USA).

GPCR functional assay was performed using the PathHunter®  $\beta$ -Arrestin assay by Eurofins (CA, USA). The PathHunter®  $\beta$ -Arrestin assay monitors the activation of a GPCR in a homogenous, non-imaging assay format using a technology called Enzyme Fragment Complementation (EFC) with  $\beta$ -galactosidase ( $\beta$ -Gal) as the functional reporter. The enzyme is split into two inactive complementary portions (EA for Enzyme Acceptor and PK for ProLink) expressed as fusion proteins in the cell. EA is fused to  $\beta$ -Arrestin and PK is fused to the GPCR of interest. When the GPCR is activated and  $\beta$ -Arrestin is recruited to the receptor, ED and EA complementation occurs, restoring  $\beta$ -Gal activity which is measured using chemiluminescent PathHunter® Detection Reagents.

The compound activity was analyzed using the CBIS data analysis suite (ChemInnovation, CA).

For agonist mode assays, percentage activity was calculated using the following formula:

%Activity = 100% x (mean RLU of test sample - mean RLU of vehicle control) / (mean MAX control ligand - mean RLU of vehicle control)

Where RLU is relative luminescence unit of the measurement.

For antagonist mode assays, percentage inhibition was calculated using the following formula:

%Inhibition = 100% x (1 - (mean RLU of test sample - mean RLU of vehicle control) / (mean RLU of EC80 control - mean RLU of vehicle control))

Where EC80 is 80% maximal effective concentration of TMAO.

### Protein-ligand docking

AutoDock Vina<sup>38</sup> was applied on TMAO to find the best conformation in the CXCR4 chemokine receptor (PDB ID: 3ODU). The center of the co-crystallized ligand (ligand ID: ITD) in 3ODU was used to define the center of the searching space and 12 Angstrom of extra space was added to the edge of ITD to set up the docking space for TMAO. The binding energies between TMAO and 3ODU were attained in terms of Kcal/mol.

### Ablation study

All the ablation studies were applied to the experiment of Microbiome-human MPI prediction.

**Vanilla TS.** Vanilla teacher-student model, where the teacher model is pre-trained, and kept frozen while training the student model, so the student will reply on the pseudo labels to learn, without sending feedback to the teacher model. Hard labels are used for pseudo-labeling.

**TS soft.** Same as **Vanilla TS**, except for soft labels are used.

**OOC-ML(out-of-cluster meta-learning).** As demonstrated in the published work<sup>12</sup>, we created five clusters based on the scaffold of the molecules, and we forced the model to see data from different clusters from every meta-update, therefore the model was pushed to generalize on the unseen data.

### Data availability

All the data used in this study can be accessed at <https://doi.org/10.5281/zenodo.10728882><sup>58</sup>; The source data can be accessed from [ChEMBL](#), [HMDB](#), and [NJS16](#); The testing dataset in Microbiome-human MPI prediction was manually created based on two published works by Chen et al.<sup>31</sup> and Colosimo et al.<sup>32</sup>. The numerical source data for all the graphs is available at: [https://figshare.com/articles/dataset/MMAPLE\\_source/26514295](https://figshare.com/articles/dataset/MMAPLE_source/26514295)<sup>59</sup>.

### Code availability

The code to reproduce results, together with documentation, is available on GitHub at <https://github.com/XieResearchGroup/MMAPLE><sup>60</sup>.

Received: 16 March 2024; Accepted: 28 August 2024;

Published online: 09 September 2024

### References

- Hicks, K. G. et al. Protein-metabolite interactomics of carbohydrate metabolism reveal regulation of lactate dehydrogenase. *Science* **379**, 996–1003 (2023).
- Chubukov, V., Gerosa, L., Kochanowski, K. & Sauer, U. Coordination of microbial metabolism. *Nat. Rev. Microbiol.* **12**, 327–340 (2014).
- Piazza, I. et al. A map of protein-metabolite interactions reveals principles of chemical communication. *Cell* **172**, 358–372 (2018).
- Groussin, M., Mazel, F. & Alm, E. J. Co-evolution and co-speciation of host-gut bacteria systems. *Cell Host Microbe* **28**, 12–22 (2020).
- Rooks, M. G. & Garrett, W. S. Gut microbiota, metabolites and host immunity. *Nat. Rev. Immunol.* **16**, 341–352 (2016).
- Bora-Tatar, G. et al. Molecular modifications on carboxylic acid derivatives as potent histone deacetylase inhibitors: Activity and docking studies. *Bioorg. Med. Chem.* **17**, 5219–5228 (2009).
- Cho, I. & Blaser, M. J. The human microbiome: at the interface of health and disease. *Nat. Rev. Genet.* **13**, 260–270 (2012).
- Xavier, J. B. et al. The cancer microbiome: distinguishing direct and indirect effects requires a systemic view. *Trends cancer* **6**, 192–204 (2020).
- Understanding the rules of life: Microbiome interactions and mechanisms (urol:mim) | nsf - national science foundation (2020). <https://new.nsf.gov/funding/opportunities/understanding-rules-life-microbiome-interactions/505694>.
- Sonnert, N. D. et al. A host–microbiota interactome reveals extensive transkingdom connectivity. *Nature* **628**, 171–179 (2024).
- Markey, K. A., van den Brink, M. R. & Peled, J. U. Therapeutics targeting the gut microbiome: rigorous pipelines for drug development. *Cell Host Microbe* **27**, 169–172 (2020).
- Cai, T. et al. End-to-end sequence-structure-function meta-learning predicts genome-wide chemical-protein interactions for dark proteins. *PLoS Comput. Biol.* **19**, e1010851 (2023).
- Sharma, K. R., Colvis, C. M., Rodgers, G. P. & Sheeley, D. M. Illuminating the druggable genome: Pathways to progress. *Drug Discovery Today* 103805 (2023).
- Kustatscher, G. et al. Understudied proteins: opportunities and challenges for functional proteomics. *Nat. Methods* **19**, 774–779 (2022).
- Oprea, T. I. et al. Unexplored therapeutic opportunities in the human genome. *Nat. Rev. Drug Discov.* **17**, 317–332 (2018).
- Xie, L., Xie, L., Kinnings, S. L. & Bourne, P. E. Novel computational approaches to polypharmacology as a means to define responses to individual drugs. *Annu. Rev. Pharmacol. Toxicol.* **52**, 361–379 (2012).
- Sadri, A. Is target-based drug discovery efficient? discovery and “off-target” mechanisms of all drugs. *J. Medicinal Chem.* **66**, 12651–12677 (2023).
- Chen, L. et al. Transformerpcpi: improving compound–protein interaction prediction by sequence-based deep learning with self-attention mechanism and label reversal experiments. *Bioinformatics* **36**, 4406–4414 (2020).
- Huang, K. et al. Deepurpose: a deep learning library for drug–target interaction prediction. *Bioinformatics* **36**, 5545–5547 (2020).
- Li, M., Lu, Z., Wu, Y. & Li, Y. Bacpi: a bi-directional attention neural network for compound–protein interaction and binding affinity prediction. *Bioinformatics* **38**, 1995–2002 (2022).
- Tossou, P., Wognum, C., Craig, M., Mary, H. & Noutahi, E. Real-world molecular out-of-distribution: Specification and investigation. *J. Chem. Inf. Modeling* **64**, 697–711 (2024).
- Liu, Y., Lim, H. & Xie, L. Exploration of chemical space with partial labeled noisy student self-training and self-supervised graph embedding. *BMC Bioinf.* **23**, 1–21 (2022).
- Lake, B. M. & Baroni, M. Human-like systematic generalization through a meta-learning neural network. *Nature* **623**, 115–121 (2023).
- Khodadadeh, S., Boloni, L. & Shah, M. Unsupervised meta-learning for few-shot image classification. *Advances in neural information processing systems* **32** (2019).
- Davies, M. et al. ChEMBL web services: streamlining access to drug discovery data and utilities. *Nucleic acids Res.* **43**, W612–W620 (2015).
- Cai, T. et al. Msa-regularized protein sequence transformer toward predicting genome-wide chemical-protein interactions: Application to gprome deorphanization. *J. Chem. Inf. Modeling* **61**, 1570–1582 (2021).
- Wishart, D. S. et al. Hmdb: the human metabolome database. *Nucleic acids Res.* **35**, D521–D526 (2007).
- Wishart, D. S. et al. Drugbank 5.0: a major update to the drugbank database for 2018. *Nucleic acids Res.* **46**, D1074–D1082 (2018).
- Liu, Y., Wu, Y., Shen, X. & Xie, L. Covid-19 multi-targeted drug repurposing using few-shot learning. *Front. Bioinf.* **1**, 693177 (2021).
- Sung, J. et al. Global metabolic interaction network of the human gut microbiota for context-specific community-scale analysis. *Nat. Commun.* **8**, 1–12 (2017).

31. Chen, H. et al. A forward chemical genetic screen reveals gut microbiota metabolites that modulate host physiology. *Cell* **177**, 1217–1231 (2019).
32. Colosimo, D. A. et al. Mapping interactions of microbial metabolites with human g-protein-coupled receptors. *Cell Host Microbe* **26**, 273–282 (2019).
33. Koeth, R. A. et al. Intestinal microbiota metabolism of l-carnitine, a nutrient in red meat, promotes atherosclerosis. *Nat. Med.* **19**, 576–585 (2013).
34. Seldin, M. M. et al. Trimethylamine n-oxide promotes vascular inflammation through signaling of mitogen-activated protein kinase and nuclear factor- $\kappa$ B. *J. Am. Heart Assoc.* **5**, e002767 (2016).
35. Yang, S. et al. Gut microbiota-dependent marker tmao in promoting cardiovascular disease: inflammation mechanism, clinical prognostic, and potential as a therapeutic target. *Front. Pharmacol.* **10**, 1360 (2019).
36. Romano, K. A., Vivas, E. I., Amador-Noguez, D. & Rey, F. E. Intestinal microbiota composition modulates choline bioavailability from diet and accumulation of the proatherogenic metabolite trimethylamine-n-oxide. *MBio* **6**, e02481–14 (2015).
37. Brunt, V. E. et al. Gut microbiome-derived metabolite trimethylamine n-oxide induces aortic stiffening and increases systolic blood pressure with aging in mice and humans. *Hypertension* **78**, 499–511 (2021).
38. Trott, O. & Olson, A. J. Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization and multithreading. *J. Comput. Chem.* **31**, 455–461 (2010).
39. Kanehisa, M. & Goto, S. Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids Res.* **28**, 27–30 (2000).
40. Yao, L. et al. Deficiency in adipocyte chemokine receptor cxcr4 exacerbates obesity and compromises thermoregulatory responses of brown adipose tissue in a mouse model of diet-induced obesity. *FASEB J.* **28**, 4534 (2014).
41. Xu, L., Kitade, H., Ni, Y. & Ota, T. Roles of chemokines and chemokine receptors in obesity-associated insulin resistance and nonalcoholic fatty liver disease. *Biomolecules* **5**, 1563–1579 (2015).
42. Costanzo-Garvey, D. L. et al. Ksr2 is an essential regulator of amp kinase, energy expenditure, and insulin sensitivity. *Cell Metab.* **10**, 366–378 (2009).
43. Liang, C.-P., Han, S., Li, G., Tabas, I. & Tall, A. R. Impaired mek signaling and serca expression promote er stress and apoptosis in insulin-resistant macrophages and are reversed by exenatide treatment. *Diabetes* **61**, 2609–2620 (2012).
44. Wong, C. K. et al. Central glucagon-like peptide 1 receptor activation inhibits toll-like receptor agonist-induced inflammation. *Cell Metab.* **36**, 130–143 (2024).
45. Diz-Chaves, Y., Mastoor, Z., Spuch, C., González-Matías, L. C. & Mallo, F. Anti-inflammatory effects of glp-1 receptor activation in the brain in neurodegenerative diseases. *Int. J. Mol. Sci.* **23**, 9583 (2022).
46. Fu, Y. et al. Gut hormone gip induces inflammation and insulin resistance in the hypothalamus. *Endocrinology* **161**, bqaa102 (2020).
47. Gray, A. L. et al.  $\alpha$ -cgrp disrupts amylin fibrillization and regulates insulin secretion: Implications on diabetes and migraine. *Chem. Sci.* **12**, 5853–5864 (2021).
48. Davis, R. B., Kechele, D. O., Blakeney, E. S., Pawlak, J. B. & Caron, K. M. Lymphatic deletion of calcitonin receptor-like receptor exacerbates intestinal inflammation. *JCI insight* **2** (2017).
49. Wei, L.-L. et al. Protective role of c3ar (c3a anaphylatoxin receptor) against atherosclerosis in atherosclerosis-prone mice. *Arteriosclerosis, Thrombosis, Vasc. Biol.* **40**, 2070–2083 (2020).
50. Finn, C., Abbeel, P. & Levine, S. Model-agnostic meta-learning for fast adaptation of deep networks. In *International conference on machine learning*, 1126–1135 (PMLR, 2017).
51. Han, B. et al. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In Bengio, S. et al. (eds.) *Advances in Neural Information Processing Systems*, vol. 31 (Curran Associates, Inc., 2018). <https://proceedings.neurips.cc/paper/2018/file/a19744e268754fb0148b017647355b7b-Paper.pdf>.
52. Thiel, C. Classification on soft labels is robust against label noise. In *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, 65–73 (Springer, 2008).
53. Uniprot: the universal protein knowledgebase in 2023. *Nucleic Acids Research* **51**, D523–D531 (2023).
54. Pham, H., Dai, Z., Xie, Q. & Le, Q. V. Meta pseudo labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11557–11568 (2021).
55. Hinton, G., Vinyals, O. & Dean, J. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015).
56. Saito, T. & Rehmsmeier, M. The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets. *PLoS ONE* **10**, e0118432 (2015).
57. Virtanen, P. et al. Scipy 1.0: fundamental algorithms for scientific computing in python. *Nat. Methods* **17**, 261–272 (2020).
58. Wu, Y. Mmaple dataset v1.0 (2024). <https://doi.org/10.5281/zenodo.10728882>.
59. Wu, Y. MMAPLE numerical source data (2024). [https://figshare.com/articles/dataset/MMAPLE\\_source/26514295](https://figshare.com/articles/dataset/MMAPLE_source/26514295).
60. Wu, Y. Xieresearchgroup/mmaple: V1.0 (2024). <https://doi.org/10.5281/zenodo.10729017>.
61. Altschul, S. F., Gish, W., Miller, W., Myers, E. W. & Lipman, D. J. Basic local alignment search tool. *J. Mol. Biol.* **215**, 403–410 (1990).

## Acknowledgements

This project has been funded with federal funds from the National Institute of General Medical Sciences of the National Institute of Health (R01GM122845), the National Institute on Aging of the National Institute of Health (R01AG057555, R21AG083302), and the National Science Foundation (2226183).

## Author contributions

Y.W. conceived the concept, prepared data, implemented the algorithms, performed the experiments, analyzed data, and wrote the manuscript; Li. X. performed the experiments, analyzed data, and wrote the manuscript; Y.L. prepared data and implemented the algorithms; Lei. X. conceived and planned the experiments, and wrote the manuscript.

## Competing interests

The authors declare that no competing interest.

## Additional information

**Supplementary information** The online version contains supplementary material available at <https://doi.org/10.1038/s42003-024-06797-z>.

**Correspondence** and requests for materials should be addressed to Lei Xie.

**Peer review information** *Communications Biology* thanks the anonymous reviewers for their contribution to the peer review of this work. Primary Handling Editors: Yuedong Yang and Laura Rodríguez Pérez. A peer review file is available.

**Reprints and permissions information** is available at <http://www.nature.com/reprints>

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Open Access** This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2024