Data Efficiency Semi-Supervised Meta-Learning Elucidates Understudied Interspecies Molecular Interactions

You Wu[1], Li Xie[2], Yang Liu[2], and Lei Xie[1,2,3,*]

[1]Ph.D. Program in Computer Science, The Graduate Center, The City University of New York, New York, New York, USA
[2]Department of Computer Science, Hunter College, The City University of New York, New York, New York, USA
[3]Helen & Robert Appel Alzheimer's Disease Research Institute, Feil Family Brain & Mind Research Institute, Weill Cornell Medicine, Cornell University, New York, New York, USA
[*]lei.xie@hunter.cuny.edu

May 17, 2023

Abstract

The power of deep learning compromises when applied to biological problems with sparsely labeled data and a data distribution shift. We developed a highly data-efficient model-agnostic semi-supervised meta-learning framework DESSML to address these challenges, and applied it to investigate understudied interspecies metabolite-protein interactions (MPI). Knowledge of interspecies MPIs is crucial to understand microbiome-host interactions. However, our understanding of interspecies MPIs is extremely poor due to experimental limitations. The paucity of experimental data also hampers the application of machine learning. DESSML successfully explores unlabeled data and transfers the information of intraspecies chemical-protein interactions to the interspecies MPI predictions. It achieves three times improvement in the prediction-recall over the baseline model. Using DESSML, we reveal novel MPIs that are validated by bioactivity assays and fill in missing links in microbiome-human interactions. DESSML is a general framework to explore previously unrecognized biological domains beyond the reach of present experimental techniques.

Key Words: machine learning, deep learning, transfer learning, pseudo-label, out-of-distribution, microbiome-host interaction, metabolite-protein interaction, protein-chemical interaction, drug-target interaction

1 Introduction

A fundamental problem in biology is to predict phenotypes under the interplay of diverse genotypes and environmental perturbations. Recent scientific advances have suggested that the microbiome co-evolves with the human and plays a critical role in shaping human phenotypes[1]. 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[2][3]. Thus, elucidating the molecular and cellular mechanisms of microbiome-host interactions will uncover the fundamental rule of life and revolutionize biomedicine[4].

Metabolites serve as not only intermediates in metabolic reactions but also signaling molecules that trigger adaptive responses through a series of molecular events including interactions with proteins in which the metabolite is not a substrate for metabolic reactions[5]. 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, such as immune response[6]. Recent proteome-wide characterization of metabolite-protein interactions (MPIs) provides strong evidence that the metabolite can play an important role in modulating gene expressions and post-modifications via off-target binding of proteins (e.g. kinases)[7]. For example, butyrate produced by the microbiome binds to human receptors such as 26S proteasome[8], HDACs[9], JMJD2E[10], and GPCRs[11]. Therefore, a comprehensive map of genome-scale interactions between microbial metabolites and human proteins will fill in a critical knowledge gap for establishing casual environment-microbiome-genotype-phenotype associations and designing small molecules to modulate microbiome-human interactions. However, our understanding of microbiome-human MPIs is limited due to the transient and low-affinity nature of interspecies MPIs[12]. In spite of tremendous advances in sequencing and high-throughput techniques, large-scale study of interspecies MPIs remains a big challenge. Existing experimental techniques for the determination of MPIs are time-consuming, biased to certain molecules, and on a relatively small scale. In this regard, *in silico* prediction of MPIs across whole microbiomes and host genomes, using machine learning, will address a crucial technical barrier in understanding microbiome–host interactions and facilitating microbiome drug discovery.

Given the significant advancement in machine learning, particularly deep learning, in recent years, it is crucial to recognize the ongoing challenges in applying this technology to understudied data-hungry problems like microbiome-human MPIs.

Due to the limitation of experimental techniques and biases of researchers' interests, the observed binding activity data is highly skewed to drug-like molecules

and a small portion of drug targets in the human genome[13][14], and reported microbiome-human MPIs are extremely scarce. Furthermore, the learning space of MPIs may differ from that of labeled drug-target interactions. Thus, the MPI prediction is a typical out-of-distribution (OOD) scenario. The data scarcity and OOD problems pose substantial obstacles to current machine learning-based approaches in predicting microbiome-human MPIs.

In this paper, we have developed a model-agnostic data efficiency semi-supervised meta-learning framework DESSML to address the challenges aforementioned. DESSML is effective in exploring unlabeled data and addressing the OOD problem. We have demonstrated that DESSML significantly improves the accuracy of microbiome-human MPI predictions by three times over state-of-the-art methods in terms of precision-recall. Using DESSML, we have identified and experimentally validated novel microbiome-human MPIs and proposed their associations with human biology. Our results suggest that DESSML can be a general framework for investigating understudied biological problems.

2 Overview of DESSML

The rationale DESSML for predicting understudied interspecies MPIs is to transfer the knowledge of intraspecies molecular interactions that are much more abundant than interspecies interactions (Figure1(A)). Specifically, we train a base model that used labeled intraspecies interactions of humans from HMDB[15] and ChEMBL[16] and microbiomes from NJS16[17]. The base model used in this study is based on a pre-trained protein language model DISAE[14] that has shown state-of-the-art performance for predicting OOD chemical-protein interactions (Figure5(A)). Then a semi-supervised meta-learning method DESSML is developed to explore unlabeled interspecies MPIs, as illustrated in Figure 1(B). DESSML first initializes a teacher model using the labeled data. Then a sampling strategy is applied to select a set of unlabeled data from the large space of understudied MPIs. The pre-trained teacher model make predictions about the selected unlabeled data and assigns labels to them (pseudo labels). Next, a student model is trained using pseudo-labeled data, at the same time, it is evaluated by labeled data and provides feedback (metadata) to the teacher model. Finally, the teacher model is updated based on the performance of the student and generates new pseudo labels. This process repeates multiple times. The details of DESSML are in the Method section.

3 DESSML significantly improves the performance of intraspecies interaction predictions

Because known interspecies MPIs are extremely scarce only including 17 observed active interactions (See Methods for details), the labeled data is insufficient to provide an unbiased assessment for DESSML. Therefore, we evaluated the performance of the DESSML model in predicting intraspecies interactions

that had a large number of available labeled data. We first trained the model using ChEMBL[16] that primarily includes exogenous small molecule ligands and druggable protein targets. We evaluated the performance of the trained model using the Human metabolite database HMDB[15] on human MPIs and NJS16[17] for microbiome MPIs. Both test cases were in the OOD setting, as supported by the protein/chemical similarity distributions (Supplemental Figure S1). The data distribution shift from ChEMBL to NJS16 is much larger than that to HMDB.

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[18]. Therefore, PR is a better metric to evaluate the performance of DESSML than ROC.

When DESSML was evaluated by HMDB, Figure2(A)indicates that DESSML significantly outperforms the state-of-the-art model DISAE on both ROC and PR. The ROC-AUC and PR-AUC increase by 20.3% and 17.4%, respectively, suggesting DESSML's ability to accurately predict human intraspecies MPIs from drug-target interactions.

Consistent with the HMDB evaluation, DESSML once again demonstrated a significant improvement over the state-of-the-art model on both ROC-AUC and PR-AUC when evaluated using NJS16, as shown in Figure2(B). There is an 11.5% and 3.8% improvement in ROC-AUC and PR-AUC, respectively. Furthermore, the trained models are not over-fitted in both cases, as supported by the narrow gaps between the training curve of validation data and that of testing data, shown in Supplemental Figure S2.

Additionally, we investigated if DESSML could alleviate the distribution shift between training and testing data. For this purpose, we extracted the embeddings of the training and testing examples acquired by DISAE and DESSML, and utilized the Uniform Manifold Approximation and Projection (UMAP) for visual analysis, and the Maximum Mean Discrepancy (MMD) for quantitative analysis. As shown in Figure2(C), DESSML moves HMDB embeddings closer to the embeddings of the ChEMBL samples globally, although the change of MMD is not large due to the overlapped embeddings in the baseline model. The effect of DESSML on NJS16 is more significant than HMDB, as shown in Figure2(D). There is almost no overlaps between the embeddings of ChEMBL samples and those of NJS16 samples. After DESSML training, two distributions have a high degree of overlap. The MMD also drops by half.

Overall, our results suggest that DESSML can be used for domain adaptation and transfer drug-target interactions to metabolite-enzyme predictions.

4 DESSML significantly improves the performance of interspecies MPI predictions

To investigate interspecies interaction, DESSML was trained on a combination of three datasets: HMDB, ChEMBL, and NJS16, while the test set consisted of 17 annotated microbiome-human MPIs along with 145 negative MPIs from the literature[19][20]. A detailed description of the datasets is available in the Method section. Our results, presented in Figure 3, demonstrate that DESSML significantly outperforms the state-of-the-art method in terms of ROC and PR. It achieves a three-fold increase in the PR for interspecies MPI predictions. These findings indicate that DESSML 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.

5 Target sampling and soft label contribute to the performance of DESSML

Our primary focus in our ablation study was on the impact of introducing target sampling, and unlabeled data, using either soft pseudo labels or hard labels.

The utilization of unlabeled data in semi-supervised learning was shown to significantly improve performance. The performance without unlabeled data and pseudo labels was suboptimal, as shown by Table 1. The removal of unlabeled data resulted in a 25% fall in ROC-AUC, and PR-AUC was less than one-third of the final outcome. Through the incorporation of unlabeled data, the model was able to explore new biological functional space effectively, resulting in better generalization.

Target sampling is a widely used technique in semi-supervised learning to increase the performance of a model by generating pseudo labels for unlabeled data. The aim of target sampling is to ensure that the generated pseudo labels have a similar distribution with the testing data, which helps the model to learn a more robust and generalizable representation of the data. When the generated pseudo labels have a different distribution than the testing data, the model may be overfitted to the pseudo-labeled data and perform poorly on the test data. As shown in Table 1, using target sampling significantly increased the performance of the model, showing the effectiveness of this strategy in improving the performance of DESSML.

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 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[21][22]. As shown in the table 1, when soft labels were used, ROC-AUC improved by 25%, and PR-AUC increased by twofold.

Table 1: Ablation study on interspecies MPIs

	ROC-AUC	PR-AUC
w/o unlabeled data	0.665 ± 0.006	0.193 ± 0.008
w/o target sampling	0.727 ± 0.033	0.308 ± 0.020
hard labels	0.708 ± 0.013	0.247 ± 0.024
DESSML	0.805 ± 0.051	0.661 ± 0.037

6 DESSML reveals the molecular basis of microbiome-human interactions

To further validate the performance of DESSML, we predicted and experimentally validated the interactions between trimethylamine N-oxide (TAMO) and human G-protein coupled receptors (CPCRs). TAMO 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[23], activate inflammatory pathways[24], and promote foam cell formation[25]. Additionally, TMAO inhibits insulin signaling[26]. However, it remains elusive how TAMO modulates these pathological processes at a molecular level. Besides its biological interest, TAMO is one of the most challenging molecules for DESSML. Firstly, the current study of microbiome-human interactions mainly focuses on short-chain fatty acids, there are few data for TAMO. Secondly, TAMO is a molecule with different structural characteristics from other chemicals in the training data. Thus, we choose TAMO to rigorously evaluate DESSML in an OOD scenario.

Figure 4(A) lists the top 7 predicted GPCR genes that interact with TAMO. We performed GPCR functional assays to experimentally test the binding activities of five of them under the concentration of 30 μM of TAMO, which is the same concentration used in the previous study and is based on the physiological concentration of TAMO in the human (1-45 μM)[27]. The assay for two top-ranked GPCRs CNRHR and ADGRA3 is not available. As shown in Figure 4(A), all five tested GPCRs show antagonist activities, and CXCR4 demonstrates a strong activity (activity score>30). The full predictive results can be found in Supplemental Table S1.

Protein-ligand docking suggested that TAMO can fit into the antagonist conformation of the CXCR4 structure, as shown in Figure 4(B) and (C). AutoDock Vina [28] was applied on TAMO to find the best conformation in the CXCR4 chemokine receptor (PDB ID: 3ODU). The docking conformations with the best (lowest) predicted binding energies were selected and the interactions between TAMO and 3ODU were shown in Figure 4(B) and (C). Among these interacting residues, TRP94, TYR116, and Glu288 also interact with the co-crystallized ligand of 3ODU. TYR116 and GLU288 provide attractive charges to the nitrogen

atom on TAMO. ARG188 forms a conventional hydrogen bond with an oxygen atom on TAMO. These strong interactions could keep TAMO in the binding pocket.

The CXCR4 antagonism by TAMO establishes a causal linkage for observed microbiome TAMO-human interactions, as illustrated in Figure 4(D). It is known that CXCR4 regulates PI3K and RAF/RAS/MEK pathways [29] (KEGG Pathway: https://www.genome.jp/pathway/hsa04062). PI3K pathway regulates bile acid synthesis[29]. TAMO's inhibition on bile acid synthesis may be responsible for its promotion effect on atherosclerosis[23]. The physiological effect of TAMO 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[30][31][32][33]. Additionally, CXCR4 is important for cell formation via the RAF/RAS/MEK pathway[34]. Thus, microbiome TAMO-human CXCR4 interaction is responsible for the several observed pathological effects of TAMO. However, several other TAMO effects such as foam cell formation and inflammation cannot be directly explained by the TAMO-CXCR4 interaction. It is possible that other human proteins can interact with TAMO.

7 Conclusion

In this study, we present DESSML, 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. Through extensive evaluations, we have demonstrated the exceptional capabilities of DESSML in exploring the unlabeled data space and facilitating knowledge transfer from one domain (intraspecies drug-protein interactions) to another (interspecies metabolite-protein interactions). Using DESSML, 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, DESSML serves as a versatile and robust framework for investigating a wide range of understudied biological problems.

DESSML shows potential for improvement in several key areas. Firstly, the current implementation of DESSML lacks the ability to estimate the uncertainty associated with pseudo labels. By incorporating an accurate uncertainty estimation 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 DESSML can be further enhanced by employing an unbiased and efficient sampling strategy. Lastly, while DESSML has thus far been applied exclusively to classification problems, it would be interesting to explore its extension to regression problems.

8 Method

8.1 Data sets

8.1.1 Protein sequence pre-training dataset

The Pfam database[35] is a comprehensive collection of protein families, represented by Multiple Sequence Alignments (MSA) and hidden Markov models (HMMs). These HMMs are used to identify proteins from new sequences that are likely to belong to a particular family, based on the presence of conserved regions. It provides a wealth of functional annotations for each protein family. It has been shown that pre-trained distilled MSA representation achieved the state-of-the-art performance[36]. We used the pre-trained model from this work directly to acquire the protein embeddings. In brief, protein sequence data were first collected from the Pfam database [35], then these sequences were then split into clusters based on the 90% of sequence identity, and a representative sequence was chosen from each cluster. The 40,282,439 sequences in the final dataset, which was used to pre-train a transformer model, included a sizable number of GPCRs. We used the original alignment as well as the conservation scores from each Pfam-A family to generate a list of amino acid triplets from the distilled sequence alignment.

8.1.2 Experiment 1: Human metabolite-protein interaction prediction

Training data The training data for this experiment was sourced from ChEMBL (version 29)[16]. It consisted of 318,687 pairs with 243,803 unique compounds and 3,208 unique proteins, where each pair represented an activity with a single protein as the target.

OOD testing data For the testing, we utilized HMDB[15], which provided interactions between metabolites and human enzymes. The dataset comprised a total of 1,156,756 pairs with 113,206 unique compounds and 3,255 unique proteins. We randomly sampled three folds, each containing 10,000 pairs. For each iteration, one fold served as the testing data. Eventually, the average score on the testing data was reported.

Unlabeled data

To create the unlabeled dataset, we considered all the unlabeled chemical-protein pairwise combinations between HMDB and ChEMBL datasets. From the total pairs, we included all unique chemicals 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.

8.1.3 Experiment 2: Microbiome metabolite-enzyme interaction prediction

Training data The training data used in this experiment was the same as that in Experiment 1.

OOD testing data

For the testing, we utilized the NJS16 dataset[17]. This dataset provided information on microbiota metabolites and their associated enzymes in microorganisms. It consisted of 162,887 pairs with 204 unique compounds and 186,284 unique enzymes. Similarly to Experiment 1, we randomly sampled three folds, each containing 10,000 pairs. One fold was used as the testing data, the process was repeated three times and the final performance was averaged.

Unlabeled data

To create the unlabeled dataset, we considered all the unlabeled chemical-protein pairwise combinations between NJS16 and ChEMBL datasets. From the full of pairs, we specifically incorporated all unique enzymes present and, in a random manner, selected two chemicals to be associated with each enzyme. This resulted in 55,476 samples in total.

8.1.4 Experiment 3: Microbiome-human MPI prediction

Training data

For this experiment, the training data consisted of a combination of ChEMBL, HMDB, and NJS16 datasets. After removing duplicates and unusable data, the dataset contained a total of 1,661,727 samples including 357,213 unique compounds and 168,517 unique proteins. To acquire the validation sets, we randomly sampled three data sets, each of which contained 10,000 samples. The remaining samples were retained for the training set.

OOD testing dataset

The testing dataset was manually created based on two published works. The first work[19] 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[20] involved the screening of gut microbiota metabolomes to identify agonists 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[37]. 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.

8.2 DESSML base model

The base model of DESSML is composed of three major modules: protein sequence modeling, chemical structure modeling, and the combination of the above two modules, as shown in Figure 5(A). The protein sequence module uses distilled sequence alignment embedding (DISAE) developed by us [14]. The protein embedding is crucial for predicting protein-ligand interactions and other biological activities in an OOD scenario [14]. DISAE uses a transformer-based architecture, which has been shown to be highly effective at learning contextual

representations of text and other sequential data. In this work, the transformer is trained on approximately half a million protein domain sequences, allowing it to learn meaningful embeddings for proteins. The chemical module is a graph neural network (GNN). This module is used to obtain chemical embeddings, which are numerical representations of small molecules and capture their chemical properties. The GNN is a type of neural network that is particularly well-suited for learning from graph-structured data, such as the chemical structures of small molecules. Finally, DESSML 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 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.

8.3 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[38]. The illustration of DESSML training is shown in Figure 5(B). Let T and S denote the teacher model and the student model, θ_T and θ_S denote the corresponding parameters (θ'_T and θ'_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.

8.4 Model training

To assure a fair comparison with the base model DISAE, both DESSML and DISAE 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

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

where

$$\mathcal{L}_u(\theta_T, \theta_S) := E_{x_u}[CE(T(x_u;\theta_T), S(x_u;\theta_S))] \tag{2}$$

The optimization of each mini-batch is performed as

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

The update rule of teacher On a batch of labeled data (x_l, y_l), and use the students' update to optimize the teacher model with the objective

$$\underset{\theta_T}{min} \mathcal{L}_l(\theta_S - \eta_S \bigtriangledown \theta_S \mathcal{L}_u(\theta_T, \theta_S)) \tag{4}$$

where

$$\mathcal{L}_l(\theta'_S) := E_{x_l, y_l}[CE(y_l, S(x_l; \theta'_S))] \tag{5}$$

The optimization of each mini-batch is performed as

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

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

Using soft labels Because we always treat θ_S as fixed parameters when optimizing Equation 6 and ignore its higher-order dependence on θ_T, the objective is fully differentiable with respect to θ_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 [39]. $T(x_u; \theta_T)$ is the teacher's output distribution computed by applying softmax over the logits z:softmax$(z) = \frac{\exp(z/T)}{\sum_{j=1}^{n} \exp(z_j/T)}$, 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[38], which was a slightly modified version of REINFORCE applied to obtain the approximated gradient of $\mathcal{L}_l$ in Equation6 with respect to θ_T as follows:

$$h = \eta_S \cdot ((\bigtriangledown_{\theta S'} CE(y_l, S(x_l; \theta_S^{(t+1)}))^T \cdot \bigtriangledown \theta_{S'} CE(\hat{y}_u, S(x_u; \theta_S^t))) \tag{7}$$

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

$$g_T^t = h \cdot \bigtriangledown \theta_T CE(\hat{y_u}, T(x_u; \theta_T))|_{\theta T} = \theta_T^{(t)} \tag{8}$$

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 θ_T with L_{base}
  end for
end for
save the model with early stopping

Stage 2:
Initialize the teacher model with θ_T
Initialize student model with random parameters θ_S
for $epoch = 1$ to n_{freeze} **do**
  for $t = 1$ to $\frac{\min(N_{un}, N_l)}{N}$ **do**
    sample $X_u n$ of size N from unlabeled data (without rep)
    update θ_S with student update rule
  end for
end for
for $epoch = n_{freeze} + 1$ to n **do**
  sample $X_u n$ of size N from unlabeled data (without rep)
  update θ_S with student update rule
  update θ_T with teacher update rule
end for

8.5 GPCR functional assay

Trimethylamine N-oxide (purity: 95%, molecular weight: 76.12) was purchased from Sigma-Aldrich (MO, USA). GPCR functional assay was performed using the PathHunter β-Arrestin assay by Eurofins (CA, USA). The compound activity was analyzed using the CBIS data analysis suite (ChemInnovation, CA).

8.6 Protein-ligand docking

AutoDock Vina [28] was applied on TAMO 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 TAMO. The binding energies between TAMO and 3ODU were attained in terms of Kcal/mol.

Acknowledgement

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), and the National Science Foundation (2226183).

Author Contributions

YW conceived the concept, prepared data, implemented the algorithms, performed the experiments, analyzed data, and wrote the manuscript; LX performed the experiments, analyzed data, and wrote the manuscript; YL conceived the concept, prepared data, implemented the algorithms, performed the experiments and analyzed data; LX conceived and planned the experiments, and wrote the manuscript.

Competing interests

The authors have declared that no competing interests exist.

References

[1] M. Groussin, F. Mazel, and E. J. Alm, "Co-evolution and co-speciation of host-gut bacteria systems," *Cell Host & Microbe*, vol. 28, no. 1, pp. 12–22, 2020.

[2] I. Cho and M. J. Blaser, "The human microbiome: at the interface of health and disease," *Nature Reviews Genetics*, vol. 13, no. 4, pp. 260–270, 2012.

[3] J. B. Xavier, V. B. Young, J. Skufca, F. Ginty, T. Testerman, A. T. Pearson, P. Macklin, A. Mitchell, I. Shmulevich, L. Xie, *et al.*, "The cancer microbiome: distinguishing direct and indirect effects requires a systemic view," *Trends in cancer*, vol. 6, no. 3, pp. 192–204, 2020.

[4] K. A. Markey, M. R. van den Brink, and J. U. Peled, "Therapeutics targeting the gut microbiome: rigorous pipelines for drug development," *Cell Host & Microbe*, vol. 27, no. 2, pp. 169–172, 2020.

[5] V. Chubukov, L. Gerosa, K. Kochanowski, and U. Sauer, "Coordination of microbial metabolism," *Nature Reviews Microbiology*, vol. 12, no. 5, pp. 327–340, 2014.

[6] M. G. Rooks and W. S. Garrett, "Gut microbiota, metabolites and host immunity," *Nature reviews immunology*, vol. 16, no. 6, pp. 341–352, 2016.

[7] I. Piazza, K. Kochanowski, V. Cappelletti, T. Fuhrer, E. Noor, U. Sauer, and P. Picotti, "A map of protein-metabolite interactions reveals principles of chemical communication," *Cell*, vol. 172, no. 1-2, pp. 358–372, 2018.

[8] G. Berkovitch, D. Doron, A. Nudelman, Z. Malik, and A. Rephaeli, "Novel multifunctional acyloxyalkyl ester prodrugs of 5-aminolevulinic acid display improved anticancer activity independent and dependent on photoactivation," *Journal of medicinal chemistry*, vol. 51, no. 23, pp. 7356–7369, 2008.

[9] C. T. Campbell, U. Aich, C. A. Weier, J. J. Wang, S. S. Choi, M. M. Wen, K. Maisel, S.-G. Sampathkumar, and K. J. Yarema, "Targeting pro-invasive oncogenes with short chain fatty acid-hexosamine analogues inhibits the mobility of metastatic mda-mb-231 breast cancer cells," *Journal of medicinal chemistry*, vol. 51, no. 24, pp. 8135–8147, 2008.

[10] D. M. Fass, R. Shah, B. Ghosh, K. Hennig, S. Norton, W.-N. Zhao, S. A. Reis, P. S. Klein, R. Mazitschek, R. L. Maglathlin, *et al.*, "Effect of inhibiting histone deacetylase with short-chain carboxylic acids and their hydroxamic acid analogs on vertebrate development and neuronal chromatin," *ACS medicinal chemistry letters*, vol. 2, no. 1, p. 39, 2010.

[11] G. Bora-Tatar, D. Dayangaç-Erden, A. S. Demir, S. Dalkara, K. Yelekçi, and H. Erdem-Yurter, "Molecular modifications on carboxylic acid derivatives as potent histone deacetylase inhibitors: Activity and docking studies," *Bioorganic & medicinal chemistry*, vol. 17, no. 14, pp. 5219–5228, 2009.

[12] J. E. Lindsley and J. Rutter, "Whence cometh the allosterome?," *Proceedings of the National Academy of Sciences*, vol. 103, no. 28, pp. 10533–10535, 2006.

[13] L. Xie and P. E. Bourne, "Functional coverage of the human genome by existing structures, structural genomics targets, and homology models," *PLOS Computational biology*, vol. 1, no. 3, p. e31, 2005.

[14] T. Cai, H. Lim, K. A. Abbu, Y. Qiu, R. Nussinov, and L. Xie, "Msa-regularized protein sequence transformer toward predicting genome-wide chemical-protein interactions: Application to gpcrome deorphanization," *Journal of chemical information and modeling*, vol. 61, no. 4, pp. 1570–1582, 2021.

[15] D. S. Wishart, D. Tzur, C. Knox, R. Eisner, A. C. Guo, N. Young, D. Cheng, K. Jewell, D. Arndt, S. Sawhney, *et al.*, "Hmdb: the human metabolome database," *Nucleic acids research*, vol. 35, no. suppl_1, pp. D521–D526, 2007.

[16] M. Davies, M. Nowotka, G. Papadatos, N. Dedman, A. Gaulton, F. Atkinson, L. Bellis, and J. P. Overington, "Chembl web services: streamlining access to drug discovery data and utilities," *Nucleic acids research*, vol. 43, no. W1, pp. W612–W620, 2015.

[17] J. Sung, S. Kim, J. J. T. Cabatbat, S. Jang, Y.-S. Jin, G. Y. Jung, N. Chia, and P.-J. Kim, "Global metabolic interaction network of the human gut microbiota for context-specific community-scale analysis," *Nature communications*, vol. 8, no. 1, pp. 1–12, 2017.

[18] T. Saito and M. Rehmsmeier, "The precision-recall plot is more informative than the roc plot when evaluating binary classifiers on imbalanced datasets," *PloS one*, vol. 10, no. 3, p. e0118432, 2015.

[19] H. Chen, P.-K. Nwe, Y. Yang, C. E. Rosen, A. A. Bielecka, M. Kuchroo, G. W. Cline, A. C. Kruse, A. M. Ring, J. M. Crawford, *et al.*, "A forward chemical genetic screen reveals gut microbiota metabolites that modulate host physiology," *Cell*, vol. 177, no. 5, pp. 1217–1231, 2019.

[20] D. A. Colosimo, J. A. Kohn, P. M. Luo, F. J. Piscotta, S. M. Han, A. J. Pickard, A. Rao, J. R. Cross, L. J. Cohen, and S. F. Brady, "Mapping interactions of microbial metabolites with human g-protein-coupled receptors," *Cell host & microbe*, vol. 26, no. 2, pp. 273–282, 2019.

[21] B. Han, Q. Yao, X. Yu, G. Niu, M. Xu, W. Hu, I. Tsang, and M. Sugiyama, "Co-teaching: Robust training of deep neural networks with extremely noisy labels," in *Advances in Neural Information Processing Systems* (S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, eds.), vol. 31, Curran Associates, Inc., 2018.

[22] C. Thiel, "Classification on soft labels is robust against label noise," in *International Conference on Knowledge-Based and Intelligent Information and Engineering Systems*, pp. 65–73, Springer, 2008.

[23] R. A. Koeth, Z. Wang, B. S. Levison, J. A. Buffa, E. Org, B. T. Sheehy, E. B. Britt, X. Fu, Y. Wu, L. Li, *et al.*, "Intestinal microbiota metabolism of l-carnitine, a nutrient in red meat, promotes atherosclerosis," *Nature medicine*, vol. 19, no. 5, pp. 576–585, 2013.

[24] M. M. Seldin, Y. Meng, H. Qi, W. Zhu, Z. Wang, S. L. Hazen, A. J. Lusis, and D. M. Shih, "Trimethylamine n-oxide promotes vascular inflammation through signaling of mitogen-activated protein kinase and nuclear factor-κb," *Journal of the American Heart Association*, vol. 5, no. 2, p. e002767, 2016.

[25] S. Yang, X. Li, F. Yang, R. Zhao, X. Pan, J. Liang, L. Tian, X. Li, L. Liu, Y. Xing, *et al.*, "Gut microbiota-dependent marker tmao in promoting cardiovascular disease: inflammation mechanism, clinical prognostic, and potential as a therapeutic target," *Frontiers in pharmacology*, vol. 10, p. 1360, 2019.

[26] K. A. Romano, E. I. Vivas, D. Amador-Noguez, and F. E. Rey, "Intestinal microbiota composition modulates choline bioavailability from diet and accumulation of the proatherogenic metabolite trimethylamine-n-oxide," *MBio*, vol. 6, no. 2, pp. e02481–14, 2015.

[27] V. E. Brunt, A. G. Casso, R. A. Gioscia-Ryan, Z. J. Sapinsley, B. P. Ziemba, Z. S. Clayton, A. E. Bazzoni, N. S. VanDongen, J. J. Richey, D. A. Hutton, *et al.*, "Gut microbiome-derived metabolite trimethylamine n-oxide induces aortic stiffening and increases systolic blood pressure with aging in mice and humans," *Hypertension*, vol. 78, no. 2, pp. 499–511, 2021.

[28] O. Trott and A. J. Olson, "Autodock vina: improving the speed and accuracy of docking with a new scoring function, efficient optimization and multithreading," *Journal of Computational Chemistry*, vol. 31, pp. 455–461, 2010.

[29] M. Kanehisa and S. Goto, "Kegg: kyoto encyclopedia of genes and genomes," *Nucleic acids research*, vol. 28, no. 1, pp. 27–30, 2000.

[30] L. Yao, J. Heuser-Baker, O. Herlea-Pana, N. Zhang, L. I. Szweda, T. M. Griffin, and J. Barlic-Dicen, "Deficiency in adipocyte chemokine receptor cxcr4 exacerbates obesity and compromises thermoregulatory responses of brown adipose tissue in a mouse model of diet-induced obesity," *The FASEB Journal*, vol. 28, no. 10, p. 4534, 2014.

[31] L. Xu, H. Kitade, Y. Ni, and T. Ota, "Roles of chemokines and chemokine receptors in obesity-associated insulin resistance and nonalcoholic fatty liver disease," *Biomolecules*, vol. 5, no. 3, pp. 1563–1579, 2015.

[32] D. L. Costanzo-Garvey, P. T. Pfluger, M. K. Dougherty, J. L. Stock, M. Boehm, O. Chaika, M. R. Fernandez, K. Fisher, R. L. Kortum, E.-G. Hong, *et al.*, "Ksr2 is an essential regulator of amp kinase, energy expenditure, and insulin sensitivity," *Cell metabolism*, vol. 10, no. 5, pp. 366–378, 2009.

[33] C.-P. Liang, S. Han, G. Li, I. Tabas, and A. R. Tall, "Impaired mek signaling and serca expression promote er stress and apoptosis in insulin-resistant macrophages and are reversed by exenatide treatment," *Diabetes*, vol. 61, no. 10, pp. 2609–2620, 2012.

[34] M. E. Bianchi and R. Mezzapelle, "The chemokine receptor cxcr4 in cell proliferation and tissue regeneration," *Frontiers in Immunology*, vol. 11, p. 2109, 2020.

[35] S. El-Gebali, J. Mistry, A. Bateman, S. R. Eddy, A. Luciani, S. C. Potter, M. Qureshi, L. J. Richardson, G. A. Salazar, A. Smart, *et al.*, "The pfam protein families database in 2019," *Nucleic acids research*, vol. 47, no. D1, pp. D427–D432, 2019.

[36] T. Cai, H. Lim, K. A. Abbu, Y. Qiu, R. Nussinov, and L. Xie, "Msa-regularized protein sequence transformer toward predicting genome-wide chemical-protein interactions: application to gpcrome deorphanization," *Journal of chemical information and modeling*, vol. 61, no. 4, pp. 1570–1582, 2021.

[37] "Uniprot: the universal protein knowledgebase in 2023," *Nucleic Acids Research*, vol. 51, no. D1, pp. D523–D531, 2023.

[38] H. Pham, Z. Dai, Q. Xie, and Q. V. Le, "Meta pseudo labels," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11557–11568, 2021.

[39] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015.

Figures

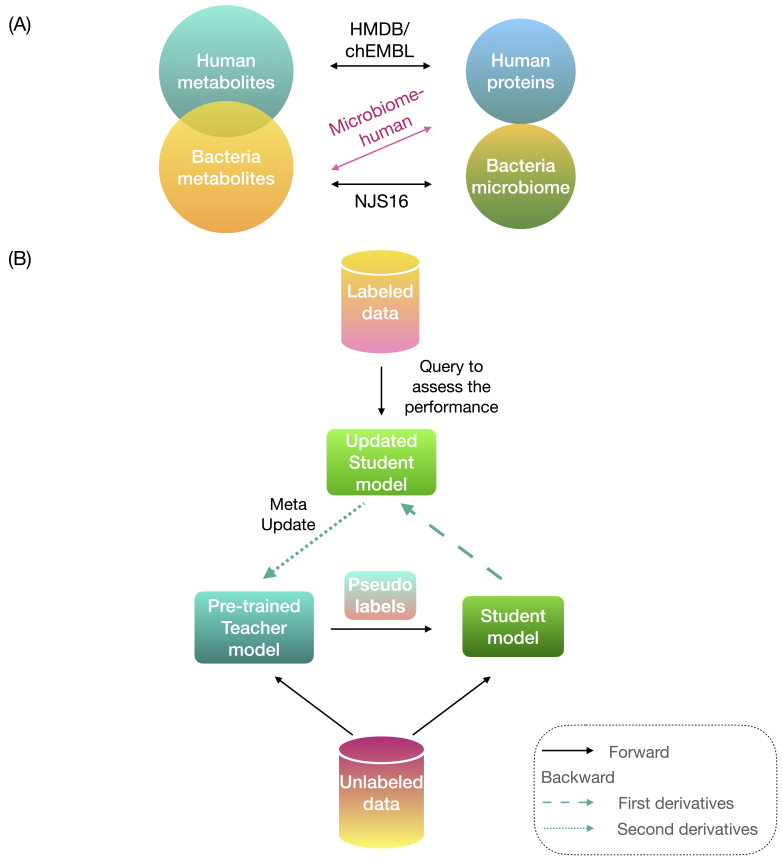

Figure 1: (A)Rationale of DESSML. The knowledge of understudied interspecies interactions is transferred from intraspecies interactions that have abundant data. (B)Illustration of DESSML framework. A deep learning model is trained using both labeled and unlabeled data and iteratively updated using gradients from the trained model as metadata.

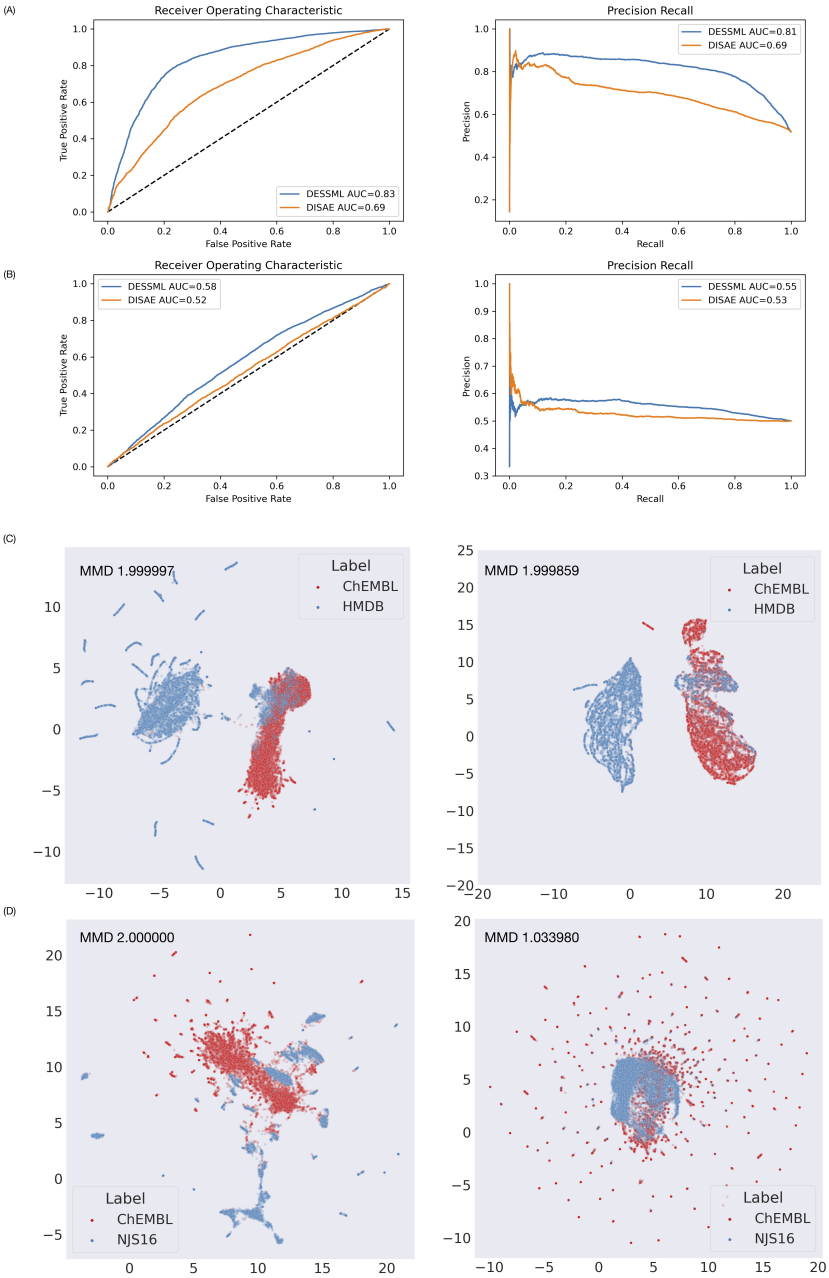

Figure 2: Performance of the Intraspecies interaction and the similarities analysis (A)Comparison between DESSML and DISAE on Human metabolite-protein association prediction (B)Comparison between DESSML and DISAE on Microbiome metabolite-enzyme association prediction; (C)Umap and MMD score of the embeddings obtained from DISAE(left) and DESSML(right) on Human metabolite-protein association on the embeddings of ChEMBL and HMDB; (D)Umap and MMD score of the embeddings obtained from DISAE(left) and DESSML(right) on Microbiome metabolite-enzyme association on the embeddings of ChEMBL and NJS16

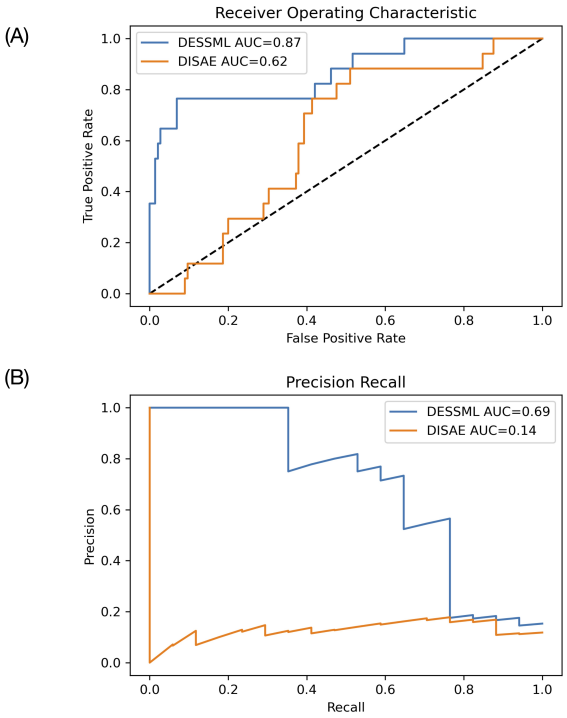

Figure 3: Performance of interspecies MPI predictions (A)ROC-AUC; (B)PR-AUC

(A)

Rank	Gene	Antagonist	Agonist	Predictive Score
1	GLP1R	3.5	0.1	0.9999683
2	GIPR	7.7	0.4	0.9999995
3	CXCR4	35.5	-0.2	0.9999988
4	CALCRL	4.1	-0.8	0.99999833
5	GNRHR	—	—	0.999998
6	ADGRA3	—	—	0.9999964
7	C3AR1	10.5	-0.4	0.99999607

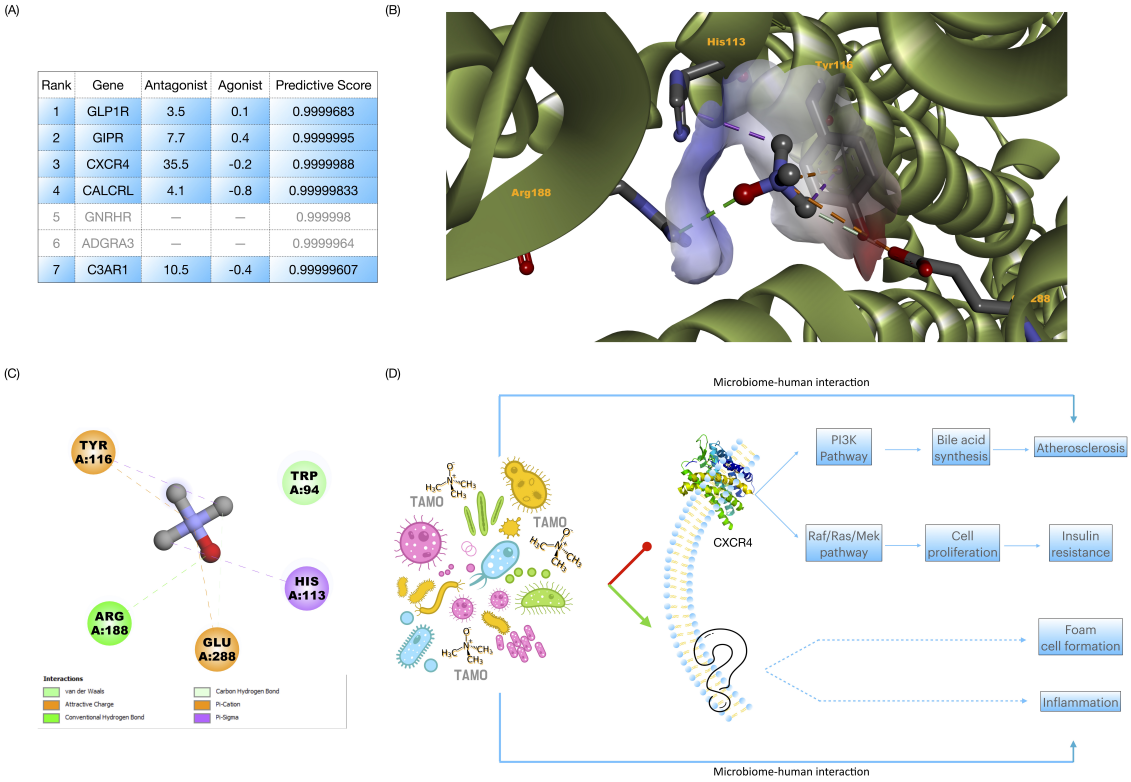

Figure 4: Experimental validation of TAMO-GPCR interactions. (A) Top 7 predicted G-protein coupled receptor (GPCR) genes that interact with TAMO and GPCR functional assay results; (B) Predicted 3D binding pose of TAMO on the CXCR4 antagonist conformation; (C) Interaction patterns between TAMO and CXCR4; (D) Proposed molecular mechanism of TAMO-human interactions. No assay is available for GNRHR and ADGRA3.

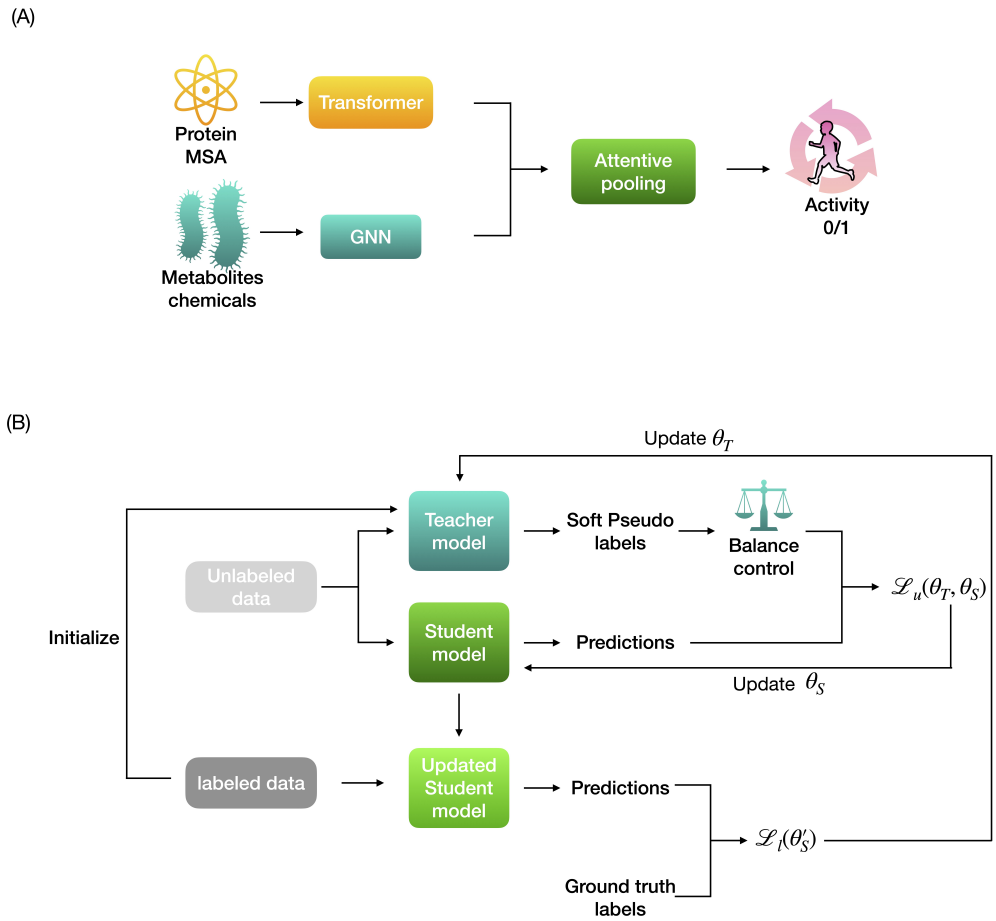

Figure 5: (A) Illustration of base neural network model for DESSML. amino acid triplets derived from protein multiple sequence alignments (MSA) are passed through a transformer architecture to generate protein embeddings; metabolite chemicals are passed through a graph neural network (GNN) to generate chemical embeddings. These two embeddings are finally combined together by attention-pooling layers to output the final result as a binary classification.
(B) Illustration of DESSML 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 hyperparameter). 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.

Supplemental materials

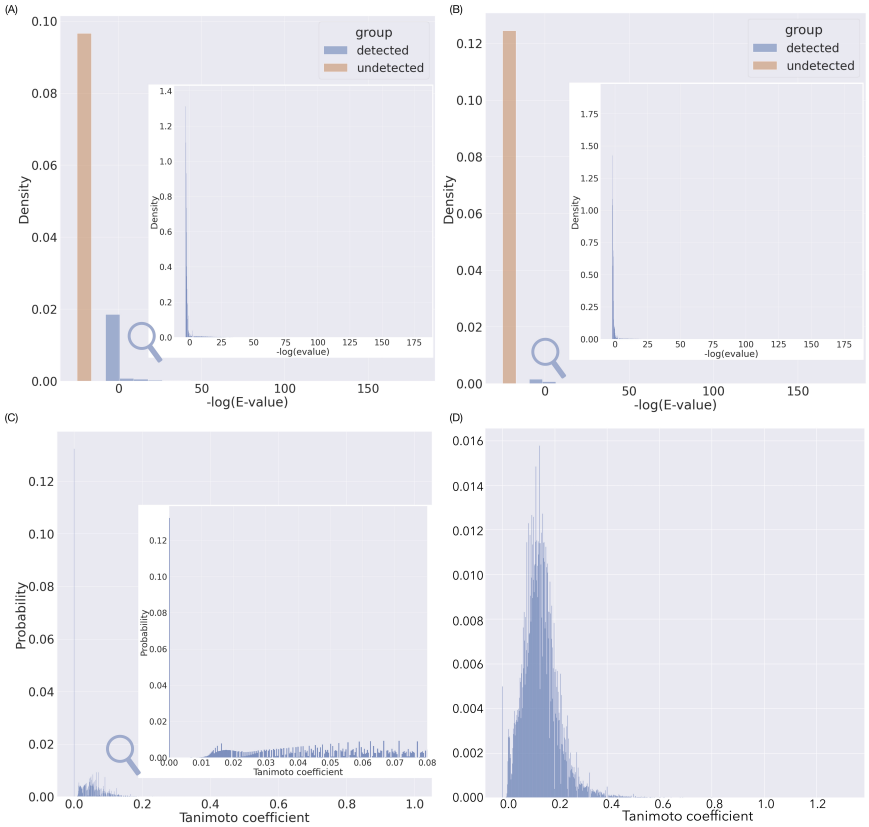

Figure S1: (A)(B)Protein similarity of Human metabolite-protein association and Microbiome metabolite-enzyme association using BLAST algorithm, undetected: E-value > 1000; (C)(D)Chemical similarity of Human metabolite-protein association and Microbiome metabolite-enzyme association measured in Tanimoto score.

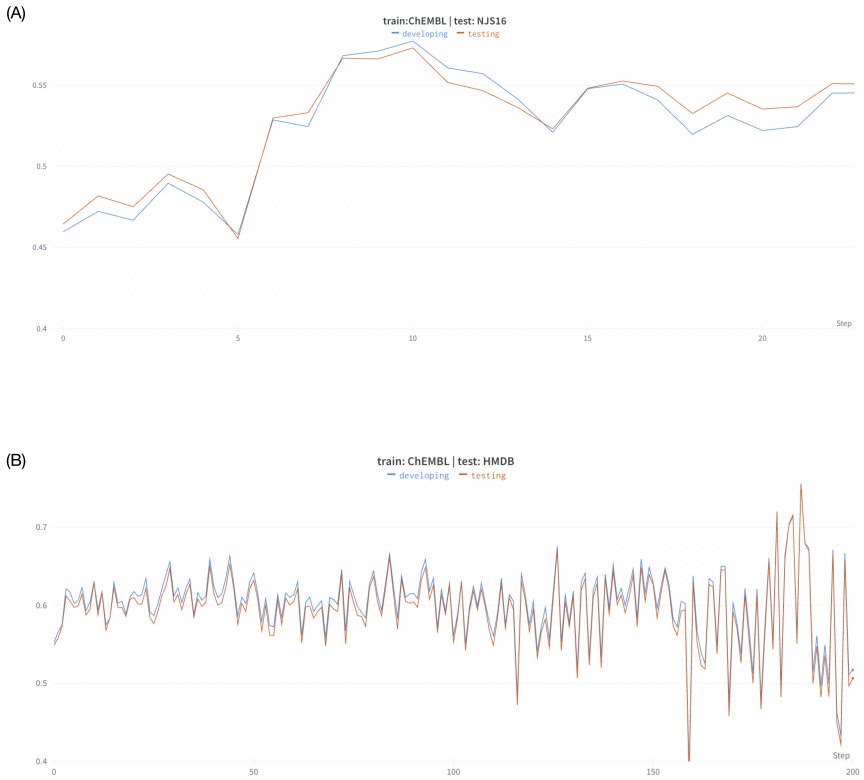

Figure S2: Learning curves of the Intraspecies interaction (A)Developing and testing curves on the Human metabolite-protein association; (B)Developing and testing curves on the Microbiome metabolite-enzyme association

Rank	Gene	Predictive score	Rank	Gene	Predictive score	Rank	Gene	Predictive score
1	GLP1R	0.9999999	32	HRH4	0.99991345	63	TAS2R14	0.9985922
2	GIPR	0.9999995	33	PTAFR	0.99990785	64	GPR151	0.99857056
3	CXCR4	0.9999988	34	FFAR4	0.9998981	65	TAS2R39	0.9981786
4	CALCRL	0.99999833	35	ADORA2B	0.9998888	66	GPR37	0.99816775
5	GNRHR	0.999998	36	KISS1R	0.99987054	67	NPY2R	0.9981553
6	ADGRA3	0.9999964	37	NPFFR1	0.9997844	68	GRM1	0.99802244
7	C3AR1	0.99999607	38	BDKRB1	0.9997336	69	NPY5R	0.9980197
8	GCGR	0.9999931	39	NPSR1	0.99971527	70	CX3CR1	0.99769956
9	CCR3	0.9999893	40	TACR2	0.9996847	71	PRLHR	0.99768734
10	GHRHR	0.9999887	41	ADGRE2	0.9996835	72	PTH1R	0.9976503
11	CCR4	0.9999871	42	HTR4	0.99961674	73	ADGRF1	0.9973182
12	CRHR1	0.99998367	43	HTR1B	0.9996137	74	ADGRD2	0.9972145
13	CALCR	0.99998355	44	ADGRG3	0.99957126	75	ADGRG2	0.9972113
14	GPR84	0.9999827	45	GRPR	0.99956053	76	TACR3	0.9969511
15	MLNR	0.99998224	46	CCR1	0.9995028	77	ADGRE3	0.9966627
16	PTGER4	0.9999795	47	ADGRB3	0.9994783	78	GALR2	0.99639326
17	NMUR2	0.99997485	48	OPRL1	0.9994055	79	GHSR	0.994869
18	SCTR	0.99996996	49	GPR143	0.9993765	80	ADRA2B	0.9944213
19	VIPR1	0.99996436	50	ADGRG7	0.99925524	81	ADGRA2	0.99412894
20	CCKAR	0.9999634	51	ACKR2	0.9992531	82	TAS2R8	0.99399406
21	MCHR1	0.99996257	52	GPR119	0.99925214	83	HRH1	0.9929675
22	CCKBR	0.9999577	53	CCR8	0.9992506	84	ADGRF3	0.9927275
23	GPR63	0.9999541	54	ADGRF4	0.99922895	85	PTGDR2	0.9923963
24	RHO	0.9999535	55	OPRM1	0.9992071	86	GLP2R	0.99078465
25	CCRL2	0.9999511	56	ADGRG4	0.9991873	87	ADRB2	0.9893908
26	GPR4	0.9999317	57	NMUR1	0.99915254	88	CYSLTR1	0.98742837
27	UTS2R	0.9999294	58	S1PR2	0.9991167	89	ADRA1A	0.9872268
28	HRH3	0.9999287	59	GPR176	0.999092	90	BDKRB2	0.9871652
29	ADORA1	0.9999199	60	ADGRG5	0.99908173	91	GPR183	0.9871394
30	HCRTR2	0.99991953	61	CXCR3	0.9990369	92	MC4R	0.9866204
31	GPBAR1	0.9999181	62	TACR1	0.9989686	93	AVPR1A	0.98517215

Table S1: Predicted top-ranked G-protein coupled receptor (GPCR) genes that interact with TAMO