

RESEARCH ARTICLE

ArcDFI: Attention regularization guided by CYP450 interactions for predicting drug-food interactions

Mogan Gim¹, Jaewoo Kang², Donghyeon Park³*, and Minji Jeon⁴*

1 Department of Biomedical Engineering, Hankuk University of Foreign Studies, Yongin, South Korea, **2** Department of Computer Science, Korea University, Seoul, South Korea, **3** Department of AI and Data Science, Sejong University, Seoul, South Korea, **4** Department of Biomedical Informatics, Korea University College of Medicine, Seoul, South Korea

* These authors are co-corresponding authors.

* parkdh@sejong.ac.kr (DP); mjjeon@korea.ac.kr (MJ)

**OPEN ACCESS**

Citation: Gim M, Kang J, Park D, Jeon M (2025) ArcDFI: Attention regularization guided by CYP450 interactions for predicting drug-food interactions. *PLOS Comput Biol* 21(10): e1013055. <https://doi.org/10.1371/journal.pcbi.1013055>

Editor: Mark Alber, University of California Riverside, UNITED STATES OF AMERICA

Received: April 14, 2025

Accepted: October 6, 2025

Published: October 15, 2025

Copyright: © 2025 Gim et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: The source code, pretrained model, and datasets are available on a GitHub repository at <https://github.com/KU-MedAI/ArcDFI>.

Funding: This work was supported by Hankuk University of Foreign Studies Research Fund (of 2024) to MG; the Korea Bio Data Station(K-BDS) with computing resources including technical support (of 2025-1) to MG; the “Regional Innovation System & Education (RISE)” through the Seoul RISE Center, funded

Abstract

CYP450 isoenzymes are known to be deeply involved in the formation of drug-food interactions (DFI). Previously introduced computational approaches for predicting DFIs do not take drug-CYP450 interactions (DCI) into account and have limited generalizability in handling compounds unseen during model training. We introduce ArcDFI, a model that utilizes attention regularization guided by CYP450 interactions to predict drug-food interactions. Experiments on DFI prediction—evaluated under stringent cold-drug and cold-food settings—show that our model outperforms ten baseline approaches, demonstrating the effectiveness of incorporating CYP450 interactions. Analysis of its attention mechanism provides insight into its current understanding of DCI and how they are related to its DFI predictions. To the best of our knowledge, ArcDFI is the first DFI prediction model that incorporates the concept of DCI, resulting in improved predictive generalizability and model explainability. ArcDFI is available at <https://github.com/KU-MedAI/ArcDFI>.

Author summary

What we eat can significantly change how medications work in our bodies. Some foods can reduce a drug’s effectiveness, while others can cause unexpected side effects. These drug-food interactions are often overlooked, yet they are critically important for patient safety. One of the main ways these interactions occur is through liver enzymes known as CYP450, which are responsible for metabolizing most clinical drugs. When foods interfere with these enzymes, they can alter how drugs are absorbed, processed, and eliminated. Despite the importance of CYP450 in drug metabolism, no previous studies have incorporated CYP450 information into computational models for predicting drug-food interactions. In this study, we present a new artificial intelligence model called ArcDFI, which is the first to explicitly include CYP450-drug interaction data

by the Ministry of Education (MOE) and the Seoul Metropolitan Government [2025-RISE-01-019-04] to DP; the National Research Foundation of Korea [RS-2024-00441029] to MJ; MSIT under the ICAN program supervised by the IITP [IITP-2025-RS-2022-00156439] to MJ; Drug Discovery R&D Center [02515431] to MJ; the Advanced Biomedical Engineering Research Group [02515711] at Korea University Anam Hospital to MJ. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

in its prediction process. Our model not only improves the ability to predict interactions involving new drugs or foods, but also explains which molecular features are most responsible for these effects. By uncovering the hidden impact of food on drug behavior, ArcDFI can support safer drug development and more personalized treatment decisions.

Introduction

Individuals commonly use medications to maintain health or treat diseases, but drug efficacy can be significantly influenced by dietary factors. Food components may alter drug absorption, metabolism, and excretion, potentially causing unexpected side effects, known as drug-food interactions (DFIs). DFIs occur via three primary mechanisms: incompatibility, pharmacokinetics (PK), or pharmacodynamics (PD) [1]. Incompatibility involves physical or chemical interactions affecting drug stability or bioavailability—e.g., tetracycline binds with calcium in dairy, reducing absorption [2]. PK interactions affect absorption, metabolism, distribution, or excretion; for example, grapefruit juice inhibits CYP3A4, increasing blood levels and toxicity risks of statins [3]. PD interactions modify drug effects at target receptors, enhancing or counteracting intended outcomes, as seen with vitamin K-rich foods diminishing warfarin's anticoagulant effect [4]. Understanding and managing DFIs is thus essential for optimizing drug therapy and ensuring patient safety.

Identifying novel DFIs is critical not only clinically but also during drug development, where early detection can improve safety profiles and accelerate approval. However, experimental identification remains difficult due to the vast number of possible drug-food pairs and complex biochemical mechanisms [5]. Therefore, computational models have emerged as scalable alternatives. DFinder, for instance, uses a graph-based method integrating compound features and topological information from a large heterogeneous graph of drugs and food [5]. DFI-MS introduces a multilevel feature optimization and contrastive learning framework [6]. While promising, both models struggle with generalization to unseen drugs or food compounds. DFinder's reliance on a fixed interaction network and DFI-MS's a compound-wise embedding layer with a fixed set of representations hinder their ability to predict interactions for novel inputs—a critical challenge for real-world applications, particularly with new therapeutics. These limitations underscore the need for more flexible, generalizable representation approaches.

Among DFI mechanisms, cytochrome P450 (CYP450)-mediated interactions are especially important, as these enzymes metabolize 75–80% of clinical drugs [7], and account for 60–70% of DFIs [8,9]. For example, grapefruit juice inhibits CYP3A4, increasing plasma levels of statins and antidepressants [10], while St. John's Wort induces it, reducing efficacy of medications such as contraceptives and immunosuppressants [11]. These examples highlight the importance of incorporating CYP450 mechanisms when studying DFIs [12].

Despite CYP450's significance, no prior computational models have explicitly used CYP450-related information for DFI prediction. This gap stems from sparse annotations for compound-CYP450 interactions. While databases such as DrugBank [13] and the Flockhart Table [14] provide some data on drug-CYP450 interactions (e.g., substrates, inhibitors), coverage remains limited, particularly for food compounds. Consequently, incorporating this critical information into existing models has been challenging. Developing sparsity-resilient models that leverage known CYP450 annotations while uncovering novel interactions is thus an important research direction.

Cross-attention mechanisms are popular in deep learning for modeling inter-modal relationships and enhancing interpretability through attention scores [15]. These scores, based

on embedding similarity between *query* and *key* elements, highlight important interactions. However, attention scores can be ambiguous when no relevant *key* exists to a given *query* element, undermining interpretability [16]. ArkDTA, a drug-target interaction model, mitigates this by using attention regularization: introducing a surrogate *pseudo*-embedding that absorbs attention when no explicit relationship exists in labeled data [17].

We extend this approach to DFI prediction using CYP450 information. Here, ground-truth labels (compound-CYP450 interactions) guide attention regularization. The *pseudo*-embedding absorbs attention if a compound lacks known interaction with the CYP450 isoenzymes, improving both prediction and interpretability. Unlike ArkDTA, however, our study faces a highly sparse set of ground truth labels for CYP450 interactions, particularly for food compounds. To address this, we apply semi-supervised learning, using attention regularization only where DCI annotations exist and defaulting to standard unsupervised cross-attention otherwise. As a result, our model can predict DFIs for novel compounds while attributing interactions to specific CYP450 isoenzymes.

In this paper, we present Attention Regularization guided by CYP450 Interactions for predicting Drug-Food Interactions (ArcDFI), a novel deep-learning framework that robustly predicts DFIs and enhances interpretability via CYP450-based attention mechanisms (Fig 1). Our key contributions are:

- Inspired by ArkDTA [17], we propose ArcDFI that employs attention regularization, which modulates its attention mechanism based on available information regarding drug-CYP450 interactions (DCIs).
- We demonstrate ArcDFI's generalizability in predicting DFI interactions given unseen drugs or food compounds based on its experimental results.
- Utilizing the attention mechanism, we demonstrate ArcDFI's explainability by finding which compound substructures are relevant in predicting DFIs and implicating which CYP450 isoenzymes are likely to form DCIs with the compound.

Materials and methods

DFI dataset construction

The main data sources for constructing the DFI dataset were FooDrugs [18] and FDMine [19]. The latest version of the FooDrugs dataset (v4) contains over 500,000 DFIs collected from textual documents using natural language processing techniques, and inferred from gene expression data using similarity profile analysis. The FDMine dataset contains binary-labeled pairwise interactions for a unique number of 787 drugs and 563 food compounds. It is a comprehensive dataset built from two large-scale data sources: DrugBank [13] and Food Database (FoodDB [20]). Note that food compounds differ from food items as the former are chemical substances, while the latter are composites of food compounds. While the data sources also contain food item information associated with food compounds, our work's main focus lies in predicting interactions between drug compounds and food compounds. Therefore, we only collected such compound-level relationship data from these data sources during the construction of our large-scale DFI dataset.

After eliminating DFIs that contain invalid compounds, we integrated both data sources to construct a large-scale DFI prediction dataset denoted as ArcDFI dataset. Positive labels indicate that a drug-food compound pair forms a metabolism-related interaction, while negative labels indicate its absence. Table 1 shows the overall statistics for the three datasets including ours.

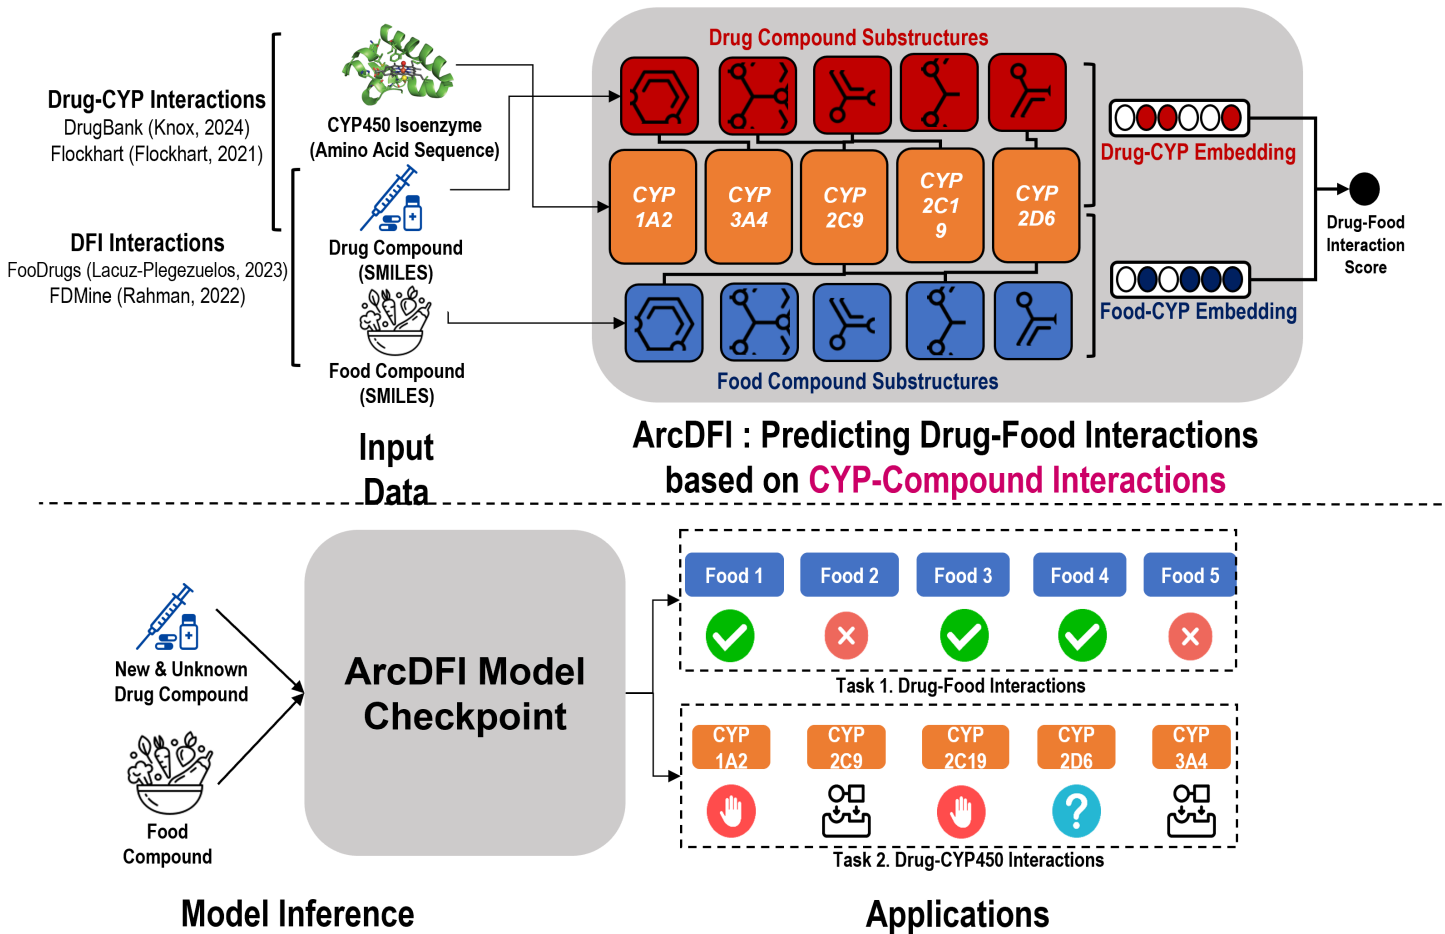


Fig 1. Overview of ArcDFI. 1) This image was obtained from https://commons.wikimedia.org/wiki/File:Capsule_icon.svg. 2) This image was obtained from <https://commons.wikimedia.org/wiki/File:Noun-drugs-1511305-00449F.svg>. 3) This image was obtained from <https://f1000research.com/articles/4-178> and is licensed under Creative Commons Attribution License, <https://doi.org/10.12688/f1000research.6314.1>.

<https://doi.org/10.1371/journal.pcbi.1013055.g001>

Table 1. Detailed statistics for each dataset where our newly constructed ArcDFI dataset is an integration of FooDrugs and FDMine.

Dataset	# of Drug Comp.	# of Food Comp.	# of DFI Pairs	# of Pos. Pairs	# of Neg. Pairs
FooDrugs (v4)	17,541	7,676	166,662	166,662	0
FDMine	786	563	442,518	2,621	439,897
ArcDFI (ours)	18,325	8,328	609,180	169,283	439,897

<https://doi.org/10.1371/journal.pcbi.1013055.t001>

Drug-CYP450 interaction label annotation

Our approach centralizes the modeling of CYP450-mediated DFIs, which requires CYP450 interaction features associated with both drug and food compounds. To augment our DFI dataset with CYP450-related information, we first gathered drug compounds annotated with CYP450 interaction labels from three data sources. The DrugBank database [13] and Drug Interactions Flockhart Table [14] provide the drug-CYP450 interactions (DCI) involving CYP450 isoenzymes and their drug interaction types. In addition, [21] released a CYP450

interaction dataset that contains both positive (presence of DCI) and negative (absence of DCI) labels for each interaction type.

After integrating the three data sources, we constructed a drug-CYP450 interaction (DCI) dataset involving five CYP450 isoenzymes (CYP1A2, CYP2C9, CYP2C19, CYP2D6, and CYP3A4), and two interaction types (substrate and inhibition). While there are other CYP450 isoenzymes (i.e., CYP2E1) and interaction types (i.e., induction), we only selected those that had relatively higher availability. The DCI dataset was used to annotate the drug compounds contained in the DFI dataset using their corresponding CYP450 interaction labels. Table 2 shows the number of drug compounds annotated for each CYP450 isoenzyme interaction type.

Model architecture for ArcDFI

As shown in Fig 2a, ArcDFI comprises the following components: Compound substructure encoder, Compound graph encoder, CYP450 encoder, Compound-CYP interaction block, Cross-modality fusion layer and Drug-food interaction prediction layer. It predicts a drug-food interaction score given both SMILES-represented compounds as input. Throughout this paper, the term "compound" is collectively used to refer to both drug compounds and food compounds.

Our architectural design aspects for implementing ArcDFI are the following,

1. **Parameter-efficient Model Design:** As DFI prediction involves two types of compound-level data entities (drug and food compounds), we incorporated weight-sharing strategies in the model architectural design of ArcDFI. Specifically, both compound types are propagated through shared Compound Encoder modules and a shared CYP-Compound Interaction block. This design reduces the number of trainable parameters and helps mitigate overfitting, particularly under limited supervision for food compound interactions.
2. **Substructure-based Representation:** We adopted ECFPs as the primary molecular representation due to their proven utility in capturing meaningful chemical substructures. ECFPs are robust, interpretable, and easily derived from canonical SMILES, making them well-suited for substructure-level attention modeling. We focused attention on substructure-enzyme interactions, as previous studies (including ArkDTA [17]) suggest these provide more interpretable and biologically meaningful insights than atom-level interactions.
3. **Attention Regularization guided by CYP450 Interactions:** To improve interpretability and inject biochemical prior knowledge, we adapted the attention regularization technique proposed in ArkDTA. While ArkDTA employed a single trainable pseudo-embedding for regularization, our method empirically benefited from using multiple pseudo-embeddings. Additionally, we divided the multi-head attention heads by

Table 2. Number of drug compounds annotated with each type of CYP450 isoenzyme (1A2, 3A4, 2C19, 2C9, 2D6) interaction (substrate, inhibition) label.

CYP450 Isoenzyme	1A2		3A4		2C19		2C9		2D6	
	Sub	Inh	Sub	Inh	Sub	Inh	Sub	Inh	Sub	Inh
FooDrugs (v4)	81	140	375	107	33	78	90	117	139	120
FDMine	92	398	329	132	57	167	78	209	170	212
ArcDFI (ours)	173	538	704	239	90	245	168	326	309	332

<https://doi.org/10.1371/journal.pcbi.1013055.t002>

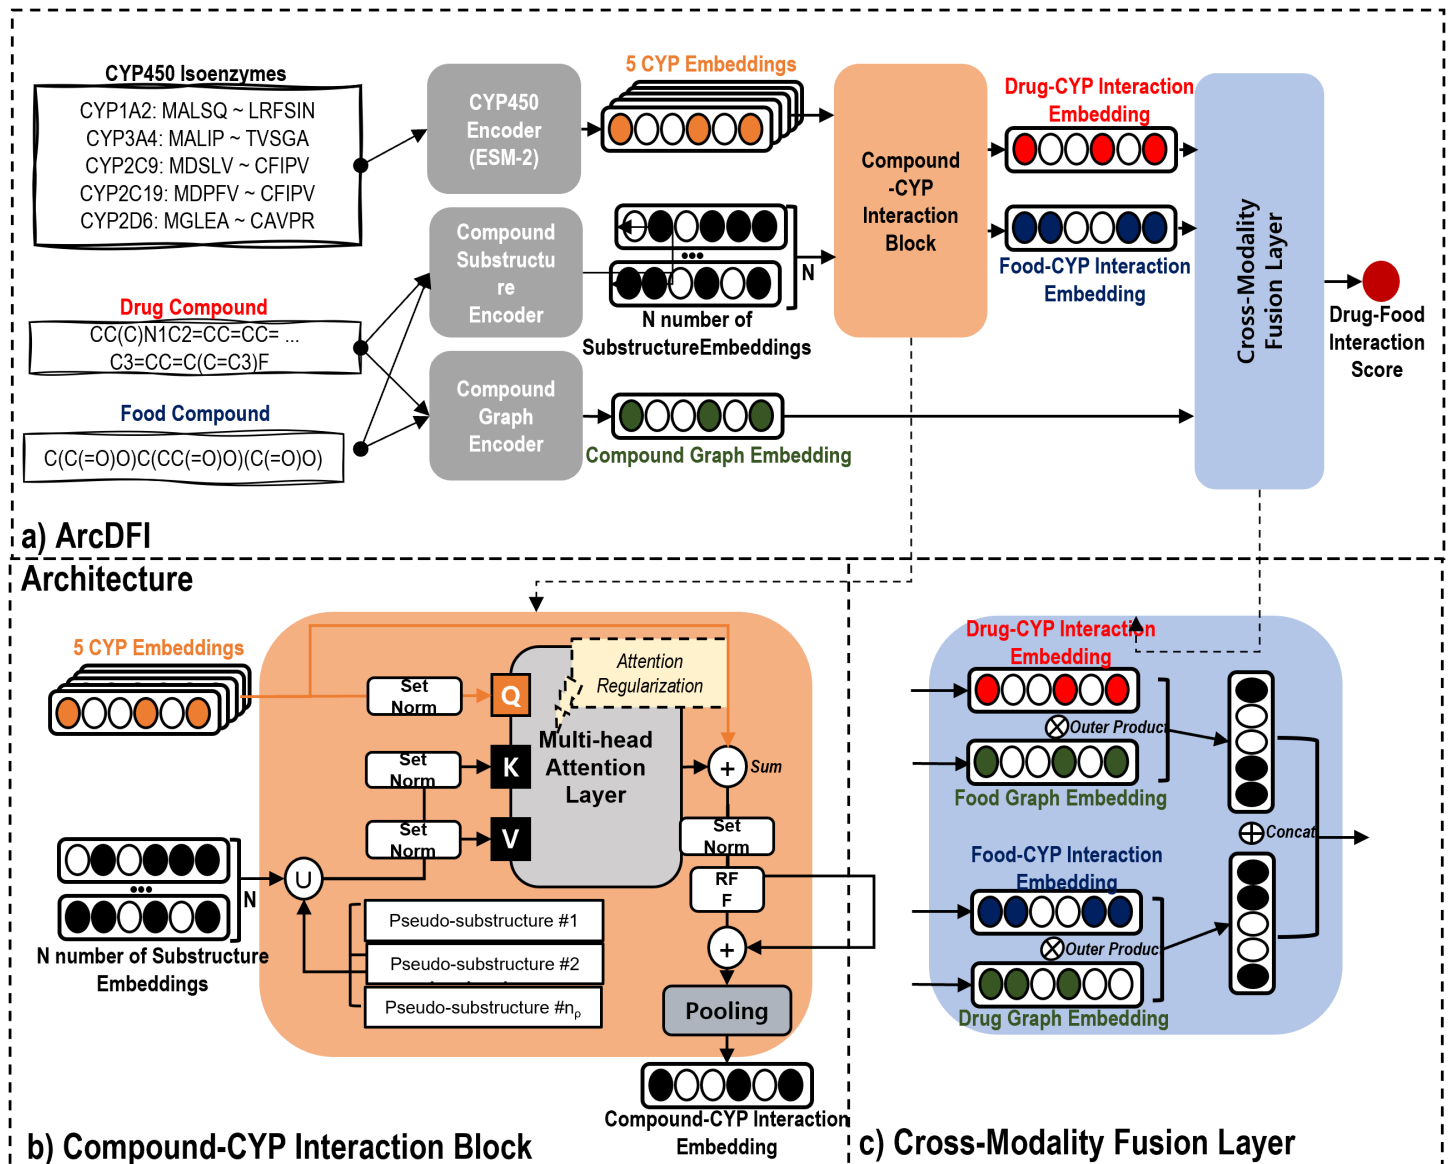


Fig 2. Descriptive illustration of ArcDFI. a) Model architecture for ArcDFI. The parameters in Compound Substructure Encoder, Compound Graph Encoder and Compound-CYP Interaction Block are shared by both drug and food compounds. b) Detailed illustration of the Compound-CYP interaction block. c) Detailed illustration of the Cross-Modality Fusion Layer. The drug (food) compound-CYP interaction embedding is combined with the food (drug) compound graph embedding using a vector-wise outer product, followed by concatenation of the two embeddings.

<https://doi.org/10.1371/journal.pcbi.1013055.g002>

the following functionalities: one head is regularized for CYP450 inhibition signals, another for substrate signals, while the remaining heads are left unregularized to retain representational flexibility.

List of mathematical notations.

For better understanding, we provide a list of mathematical notations base below.

- d* Drug Compound
- f* Food Compound

\mathcal{S}	Chemical Substructures of a Compound
\mathcal{G}	Molecular Graph Representation of a Compound
\mathcal{E}	CYP450 IsoEnzymes
ρ	Trainable Pseudo-substructures
$\mathbf{X}^{\mathcal{S}}$	Initial Set of \mathcal{S} Embeddings
$\mathbf{X}^{\mathcal{G}}$	Initial Compound Graph Representation
$\mathbf{X}^{\mathcal{E}}$	Set of 5 Amino-Acid Sequences of \mathcal{E}
$\mathbf{H}^{\mathcal{S}}$	Set of Encoded Chemical Substructure Embeddings for a Compound
$\mathbf{H}^{\mathcal{G}}$	Molecular Graph Embedding for a Compound
$\mathbf{H}^{\mathcal{E}}$	Set of 5 Refined CYP450 Isoenzyme Embeddings
\mathbf{H}^{ρ}	Set of Pseudo-Substructure Embeddings
$\mathbf{H}_d^{\mathcal{I}}$	Drug Compound-Integrated CYP450 Embedding
$\mathbf{H}_f^{\mathcal{I}}$	Food Compound-Integrated CYP450 Embedding
$\mathbf{H}^{\mathcal{F}}$	CYP450-mediated Drug-Food Fusion Embedding
$\hat{\mathbf{Y}}$	Batch of DFI Score Predictions
\mathbf{Y}	Batch of Ground Truth DFI Labels
$\mathcal{A}_{i,j,k}$	Attention Weight Score calculated between the j th CYP450 Isoenzyme and k th Compound Substructure in the i th Attention Head
$\hat{\mathbf{Y}}_c$	Batch of CYP450 Interaction Score Predictions
\mathbf{Y}	Batch of Ground Truth CYP450 Interaction Labels

Compound substructure encoder. The Compound substructure encoder first takes a SMILES- represented compound as input and converts it into 1024-dimensional Extended Connectivity Fingerprints (ECFPs) with a radius set to 2. Because each bit position of its binary ECFP represents the presence of its substructure detected within a certain radius, we utilized this information by converting each compound into a set of $2h$ -dimensional trainable substructure embeddings [17]. These initial substructure embeddings corresponding to the drug and food compounds were then propagated through an MLP, resulting in a set of encoded n_d and n_f chemical substructure embeddings, respectively, where d and f represent a drug compound and a food compound respectively ($\mathbf{X}_d^{\mathcal{S}} \in \mathbb{R}^{n_d \times h}$, $\mathbf{X}_f^{\mathcal{S}} \in \mathbb{R}^{n_f \times h}$).

The Compound substructure encoder that takes the ECFPs from SMILES representation of drugs and food compounds as inputs is mathematically expressed as follows:

$$\mathbf{H}^{\mathcal{S}} = \text{MLP}(\mathbf{X}^{\mathcal{S}}) \quad (1)$$

$$\text{MLP}(x) = \text{Dropout}(\text{LeakyReLU}(xW + b)) \quad (2)$$

where $\mathbf{X}^{\mathcal{S}} \in \mathbb{R}^{n \times 2h}$ is an initial set of $2h$ -dimensional chemical substructure embeddings converted from the compound's ECFP representation. MLP comprises a linear transformation layer ($W \in \mathbb{R}^{2h \times h}$, $b \in \mathbb{R}^h$) followed by non-linear activation function LeakyReLU and Dropout layer whose dropout rate is globally set to 0.3.

Compound graph encoder. The Compound graph encoder also uses a SMILES- represented compound as input and converts it into its molecular graph representation ($\mathbf{X}^{\mathcal{G}}$) compatible with its inherent Graph Isomorphism Network (GIN) convolution layer that additionally incorporates edge attributes [22]. The initial features for the nodes (atoms) are the atomic number, chirality, degree, formal charge, number of hydrogen atoms, number of radical electrons, hybridization, aromaticity, and ring-like structure. The initial features of the edges (bonds) are the bond types, stereo configuration, and conjugation.

Atom-wise node embeddings of the given compound were built using the GIN convolution layer based on the initial atom-bond information and topological characteristics. They are subsequently aggregated through a global mean pooling layer into an h -dimensional single-handed molecular graph embedding which is mathematically expressed as $\mathbf{H}_d^G \in \mathbb{R}^h$ and $\mathbf{H}_f^G \in \mathbb{R}^h$ respectively.

The Compound graph encoder, which uses the SMILES representation of the drug and food compounds as input is mathematically expressed as follows:

$$\mathbf{N}^G = \text{GINEConv}(\mathbf{X}^G) \quad (3)$$

$$\mathbf{H}^G = \text{MeanPool}(\mathbf{N}^G) \quad (4)$$

where \mathbf{X}^G is the molecular graph of the input compound previously converted from its SMILES representation and contains the initial atom and bond features. GINEConv is the GIN convolution layer using edge attribute features, and MeanPool is a global mean pooling layer that aggregates the node embeddings \mathbf{N}^G to a single molecular graph embedding $\mathbf{H}^G \in \mathbb{R}^h$.

CYP450 encoder. The CYP450 encoder takes a set of five CYP450 isoenzymes initially represented as amino acid sequences as inputs and converts them into a set of contextualized CYP450 isoenzyme embeddings. The input amino acid sequence representations for the isoenzymes were obtained from the UniProt database [23]. They were fed into ESM-2, a large-scale protein language transformer model trained on millions of protein sequences [24], which resulted in a set of five 480-dimensional language-based embeddings $\mathbf{L}^E \in \mathbb{R}^{5 \times 480}$. Lastly, these language-based embeddings are propagated through an MLP, which results in a set of 5 h -dimensional refined protein embeddings $\mathbf{H}^E \in \mathbb{R}^{5 \times h}$. Note that the contextualized CYP450 isoenzyme embeddings are universally used for all drug-food compound pairs within the same batch during training.

The CYP450 encoder, which takes the five CYP450 isoenzymes \mathbf{X}^E , each represented as an input sequence of amino acids, is mathematically expressed as follows:

$$\mathbf{L}^E = \text{ESM-2}(\mathbf{X}^E) \quad (5)$$

$$\mathbf{H}^E = \text{MLP}(\mathbf{L}^E) \quad (6)$$

$$\text{MLP}(x) = \text{Dropout}(\text{LeakyReLU}(xW^T + b)) \quad (7)$$

where $\mathbf{L}^E \in \mathbb{R}^{5 \times 480}$ are the language-based embeddings for the five CYP450 isoenzymes encoded by the protein language transformer model, ESM-2 (esm2-t12-35M-UR50D). MLP comprises a linear transformation layer ($W \in \mathbb{R}^{480 \times h}$, $b \in \mathbb{R}^h$) followed by LeakyReLU and a Dropout layer.

Compound-CYP interaction block. The Compound-CYP interaction block (Fig 2b) calculates multi-head cross-attention scores between the contextualized CYP450 isoenzyme and chemical substructure embeddings, which are the input queries and key values, respectively. Its design intuition aligns with a ligand (food or drug compound) binding to the target CYP450 isoenzyme, exhibiting either inhibitory or substrate-related effects. By interpreting pairwise attention scores calculated by the cross-attention mechanism, we can imply which substructures play a crucial role in DCIs.

The cross-attention mechanism usually involves the calculation of pairwise attention weights between CYP450 isoenzymes (queries) and chemical substructure embeddings (keys) in an unsupervised manner. Attention weights are distributed to the keys for each query, regardless of whether the compound actually interacts with the isoenzyme. To address this

issue, we adopted ArkDTA's attention regularization method, which aims to modulate the distribution of attention weights driven by an auxiliary loss objective [17].

Specifically, a set of n_ρ trainable h -dimensional pseudo-substructure embeddings are first appended to the current set of chemical substructure embeddings. Attention weights distributed to the pseudo-substructures imply absence of DCIs while the opposite applies for those distributed to the actual ones. We treated the sum of the weights assigned to both the pseudo-substructures and the actual substructures as binary class probability scores. They are fed to a binary cross-entropy loss function, where the ground truth labels are the compound-CYP450 interaction annotations from our DCI dataset.

Given that our study included two types of DCI substrates and inhibition, we applied supervised attention regularization to the attention weights from the first and second heads of the block, corresponding to substrate and inhibition effects, respectively. The remaining attention heads were unsupervised during the training process. The output of the Compound-CYP interaction block is an aggregation of 5 h -dimensional attention-based CYP450 isoenzyme embeddings $\mathbf{O}^\mathcal{E} \in \mathbb{R}^{5 \times h}$ using mean pooling. We denote this final output as compound-integrated CYP450 embedding $\mathbf{H}^\mathcal{I} \in \mathbb{R}^h$. A detailed explanation of the loss objective is available in the Auxiliary Loss Objective section.

The Compound-CYP interaction block that takes the set of five contextualized CYP450 isoenzymes $\mathbf{H}^\mathcal{E}$ and n chemical substructure embeddings $\mathbf{H}^\mathcal{S}$ as inputs is mathematically expressed as follows:

$$\mathbf{H}^{\mathcal{S}+\rho} = \mathbf{H}^\mathcal{S} \cup \mathbf{H}^\rho \quad (8)$$

$$\mathbf{O}^\mathcal{E} = \text{MAB}(\mathbf{H}^\mathcal{E}, \mathbf{H}^{\mathcal{S}+\rho}) \quad (9)$$

$$\mathbf{H}^\mathcal{I} = \text{MeanPool}(\mathbf{O}^\mathcal{E}) \quad (10)$$

$$\text{MAB}(x, y) = (x + z) + \text{RFF}(\text{SetNorm}(x + z)) \quad (11)$$

$$z = \text{Multihead}(\text{SetNorm}(x), \text{SetNorm}(y), \text{SetNorm}(y)) \quad (12)$$

$$\text{RFF}(x) = \text{GeLU}(\text{GeLU}(xW_1 + b_1)W_2 + b_2) \quad (13)$$

where MAB is a multihead attention block borrowed from the Set Transformer architecture [25] enhanced by set normalization layers and equivariant skip connections [26]. The attention block includes a multihead attention layer, where the input query, key, and value embeddings are $\mathbf{H}^\mathcal{E}$, $\mathbf{H}^{\mathcal{S}+\rho}$ and $\mathbf{H}^{\mathcal{S}+\rho}$ respectively. It employs four attention heads with pairwise attention weights computed using additive attention. Note that the computed attention weights from the first two heads are fed to an auxiliary loss objective. RFF is a row-wise feed-forward layer consisting of two layers of linear transformation ($W_1, W_2 \in \mathbb{R}^{h \times h}$, $b_1, b_2 \in \mathbb{R}^h$) each followed by Gaussian Error Linear Unit (GeLU) activation function.

Cross-modality fusion layer. To effectively model the complex interactions between CYP450 isoenzymes, drug compounds, and food compounds, we implemented a Cross-modality fusion layer (Fig 2c) that combines attention-based drug-CYP450 integrative features $\mathbf{H}_d^\mathcal{I}$ with the food compound's graph topological features $\mathbf{H}_f^\mathcal{G}$ and vice versa ($\mathbf{H}_f^\mathcal{I}$, $\mathbf{H}_d^\mathcal{G}$). Recall that $\mathbf{H}_d^\mathcal{G}$ and $\mathbf{H}_f^\mathcal{G}$ represent molecular graph embeddings for the drug and food compound, respectively, whereas $\mathbf{H}_d^\mathcal{I}$, $\mathbf{H}_f^\mathcal{I}$ denote the CYP450 embeddings integrated with the drug and food compound, respectively.

The fusion process employs a vector-wise outer product between the two heterogeneous embeddings for both sides of the drug-food pair. Each of the outer product results is reshaped

into a single h^2 -dimensional embedding. Lastly, the two reshaped embeddings are concatenated to each other vector-wise, producing a final output of this layer: a $2h^2$ -dimensional CYP450-mediated drug-food fusion embedding $\mathbf{H}^{\mathcal{F}} \in \mathbb{R}^{2h^2}$.

The Cross-modality fusion layer that takes the compound-integrated CYP450 isoenzyme embeddings $\mathbf{H}_d^{\mathcal{I}}, \mathbf{H}_f^{\mathcal{I}}$ and molecular graph embeddings $\mathbf{H}_f^{\mathcal{G}}, \mathbf{H}_d^{\mathcal{G}}$ as input is mathematically expressed as follows,

$$\mathbf{H}^{\mathcal{F}} = \text{Reshape}(\mathbf{H}_d^{\mathcal{I}} \otimes \mathbf{H}_f^{\mathcal{G}}) \oplus \text{Reshape}(\mathbf{H}_f^{\mathcal{I}} \otimes \mathbf{H}_d^{\mathcal{G}}) \quad (14)$$

where Reshape, \otimes , \oplus refers to the reshaping process from $\mathbb{R}^{h \times h}$ matrix to \mathbb{R}^{h^2} vector, vector-wise outer product and concatenation.

Drug-food interaction prediction layer. Finally, the Drug-food interaction prediction layer takes the single CYP450-mediated drug-food fusion embedding as input and ultimately predicts the interaction likelihood between the drug and food compound. The Drug-food interaction prediction layer is a deeply stacked feedforward layer that takes the single CYP450-mediated drug-food fusion embedding as input and ultimately predicts the interaction likelihood between the drug and food compound. This can be mathematically expressed as follows:

$$\mathbf{H}_1 = \text{Dropout}(\text{LeakyReLU}(\text{BatchNorm}(\mathbf{H}^{\mathcal{F}} W_1 + b_1))) \quad (15)$$

$$\mathbf{H}_2 = \text{Dropout}(\text{LeakyReLU}(\mathbf{H}_1 W_2 + b_2)) \quad (16)$$

$$\mathbf{H}_3 = \text{Dropout}(\text{LeakyReLU}(\text{BatchNorm}(\mathbf{H}_2 W_3 + b_3))) \quad (17)$$

$$\hat{y} = \text{Sigmoid}(\mathbf{H}_3 W_4 + b_4) \quad (18)$$

where BatchNorm and Sigmoid refer to the batch normalization and sigmoid functions, respectively. Weights and bias for the linear transformation layers are $W_1 \in \mathbb{R}^{2h^2 \times \lceil \sqrt{2h^2} \rceil}$, $W_2 \in \mathbb{R}^{\lceil \sqrt{2h^2} \rceil \times 2h}$, $W_3 \in \mathbb{R}^{2h \times h}$, $W_4 \in \mathbb{R}^{h \times 1}$ and $b_1 \in \mathbb{R}^{\lceil \sqrt{2h^2} \rceil}$, $b_2 \in \mathbb{R}^{2h}$, $b_3 \in \mathbb{R}^h$, $b_4 \in \mathbb{R}^1$ respectively.

Model optimization

ArcDFI is trained under two loss objectives which are the primary loss objective for predicting DFIs and auxiliary loss objective for attention regularization. The former is implemented under supervised learning setting.

Primary loss objective. The batch-wise primary loss objective for drug-food interaction prediction, treated as a binary classification task, is mathematically expressed as follows:

$$\mathcal{L}_1(\hat{\mathbf{Y}}, \mathbf{Y}) = \text{BinaryCrossEntropy}(\hat{\mathbf{Y}}, \mathbf{Y}) \quad (19)$$

where $\hat{\mathbf{Y}} \in \mathbb{R}^b$ and $\mathbf{Y} \in \{0, 1\}^b$ is a batch of predicted DFI scores and ground truth binary labels respectively, and b is the batch size.

Auxiliary loss objective. To enhance ArcDFI's understanding of CYP450-mediated drug-food interactions through semi-supervised learning, we introduced an auxiliary loss objective to regularize the attention mechanism. Let $\mathbb{A}_{i,j,k}$ denote the attention weights computed between the j th CYP450 isoenzyme and k th compound substructure in the i th attention head within the Compound-CYP interaction block. The value of $\mathbb{A}_{i,j,k}$ can be interpreted as the likelihood of the k th compound substructure contributing to the substrate ($i = 1$) or inhibitory effects ($i = 2$) to the j th CYP450 isoenzyme. To clarify, we applied attention regularization to

the first ($\mathbb{A}_{1,j,k}$) and second attention heads ($\mathbb{A}_{2,j,k}$), whereas the remainder were left unsupervised as originally designed.

As described earlier, the Compound-CYP interaction block module appends a universal set of n_ρ trainable pseudo-substructure embeddings to the current set of n key substructure embeddings resulting in a total of $n + n_\rho$ embeddings involved in pairwise attention weight computation with respect to the five query CYP450 isoenzyme embeddings. Given $\mathbb{A}_{1,j,k}$, if the partner compound forms a substrate binding interaction with the j th CYP450 isoenzyme according to the ground truth label in our DCI dataset, our designed auxiliary loss objective encourages the distribution of attention weights to concentrate on the n actual substructure embeddings $\mathbb{A}_{1,j,1:n}$, while reducing the weights for the n_ρ pseudo-substructure embeddings, $\mathbb{A}_{1,j,n+1:n+n_\rho}$. Similarly, this modulation applies to $\mathbb{A}_{2,j,k}$ for compounds exhibiting inhibitory effects on the j th CYP450 isoenzyme.

Conversely, when the ground truth label for a specific DCI type is negative, the loss objective promotes the opposite behavior by diverting attention away from the actual substructure embeddings. If the DCI label is unavailable, the loss objective is not applied, leaving the attention distribution unsupervised. Note that the calculated pairwise attention weights for each row in $\mathbb{A}_{i,j,k}$ are normalized using a softmax function, making them inherently suitable for optimization using a cross-entropy loss in the auxiliary objective.

The batch-wise auxiliary loss objective for attention regularization based on compound-CYP450 interactions is mathematically expressed as follows:

$$\mathcal{L}_2(\mathbb{A}, \mathbf{Y}_c) = \text{BinaryCrossEntropy}(\hat{\mathbf{Y}}_c, \mathbf{Y}_c) \quad (20)$$

where $\hat{\mathbf{Y}}_c \in \mathbb{R}^{b \times 10}$ and $\mathbf{Y}_c \in \{0, 1\}^{b \times 10}$ is a batch of predicted interaction scores and ground truth binary labels respectively for each CYP450 isoenzyme and type. The interaction scores $\hat{\mathbf{Y}}_c$ are calculated based on the computed attention weights extracted from the Compound-CYP interaction block.

Note that the sparse proportion of drug compounds with CYP450 isoenzyme interaction labels sets the auxiliary loss objective in a semi-supervised learning setting. When interaction labels are unavailable, the corresponding predicted interaction scores remain unsupervised. Only the drug compounds with available interaction labels are used to update the parameters of the Compound-CYP interaction block. Fig 3 provides an illustrated description for the auxiliary loss objective.

Training ArcDFI. Conclusively, the batch-wise total loss objective for training ArcDFI is mathematically expressed as follows,

$$\mathcal{L} = \mathcal{L}_1 + \alpha \mathcal{L}_2 \quad (21)$$

where α is the auxiliary loss coefficient for controlling the intensity of attention regularization based on CYP450 interaction.

We trained ArcDFI for a maximum of 100 epochs using early stopping, optimized with the AdamW optimizer. The hyperparameters for ArcDFI's model configuration include batch size, learning rate, weight decay, auxiliary loss coefficient, hidden embedding dimensionality, and the number of pseudo-embeddings. The batch size (b), learning rate, and weight decay were set to 1024, 0.0001, and 0.0001, respectively. Early stopping was configured with a patience of 10 epochs, using the validation loss as the monitoring metric. The auxiliary loss coefficient α was set to 5. All hyperparameters for the model architecture and training algorithm were consistently the same across all settings involving ArcDFI. For the model-related

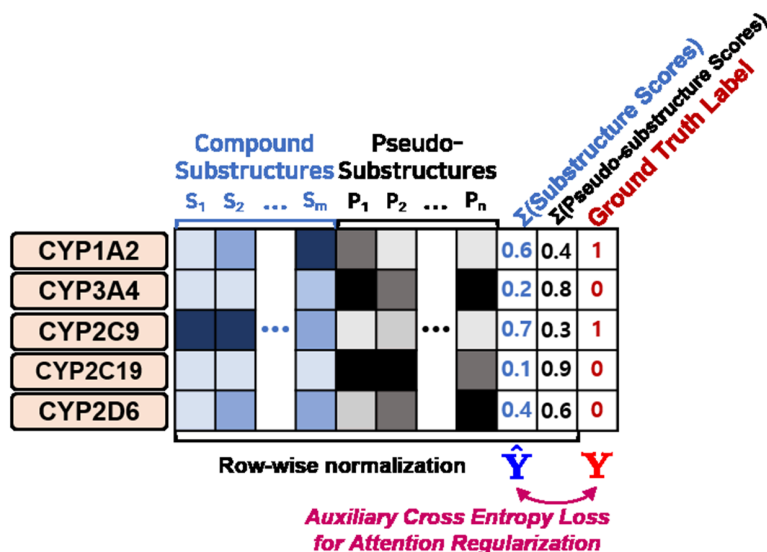


Fig 3. Detailed illustration for Attention Regularization Auxiliary Loss Objective.

<https://doi.org/10.1371/journal.pcbi.1013055.g003>

configurations of ArcDFI, the hidden dimensional size of embeddings and the number of pseudo-substructure embeddings were set to $h = 128$ and $\rho = 10$. We initially referenced the hyperparameter settings from the ArkDTA paper and then further refined them through a series of experiments, optimizing based on our model's performance on the validation set.

Results

Evaluation on DFI prediction

To evaluate ArcDFI's performance in DFI prediction, we first selected DFinder [5] and DFI-MS [6] as our main baselines. These models were originally designed to predict DFIs when given a pair of drug and food compounds as input. Since the DFI prediction task is an under-explored topic compared to DDI, we decided to include more baselines that are not necessarily tailored for DFI prediction, but compatible with the task itself, as this prediction task involves building learnable representations from input compounds. The following baseline models originally designed for DDI prediction were included in our DFI prediction experiments: DeepSynergy [27], DeepDDI [54], EPGCN-DS [29], CASTER [30], SSI-DDI [31], MatchMaker [32], MR-GNN [33], and DeepDrug [28].

Unlike previous studies on DFI, we employed a stricter evaluation approach that involved splitting the dataset based on its compound clusters. We first utilized Butina clustering to build drug and food compound clusters in the DFI dataset [34] and then made two data splits based on the food or drug compound clusters, which are referred to as cold drug and cold food settings. This data split approach has been widely used in research on drug-target interaction prediction [17,35,36]. The resultant number of drug-food compound pair-based data instances split to training, validation, and test partitions in each cold setting is the following,

Evaluation Setting	Training	Validation	Test
Cold-Drug Split	444,018	33,252	137,735
Cold-Food Split	458,827	23,174	133,004

We repeated the experiments three times using different random seeds for each model and data-split setting to ensure stability and robustness. We used five evaluation metrics to measure the predictive performance: Area Under ROC Curve (AUROC), Area Under Precision-Recall Curve (AUPRC), F1-score (F1), Precision, and Recall. The final scores for each model were averaged across three runs for each split setting.

During the DFI experiments, models, including baselines, were trained and evaluated on the same DFI dataset with different splits according to the evaluation setting (cold-drug, cold-food). Also, all baseline models imported from their respective GitHub repositories were trained with their default hyperparameters.

Tables 3 and 4 show the quantitative results of evaluating the new drug and food setting, respectively. ArcDFI outperformed its baseline models for the cold drug setting, especially in terms of AUROC, AUPRC, and F1, and even its ablated version. While cold food evaluation setting presented a harder challenge to the DFI models, ArcDFI showed second-best predictive performance. In contrast, the ablated version of ArcDFI, in which its auxiliary loss objective was removed, outperformed other models in the cold food setting. These results demonstrated that the proposed attention regularization method is advantageous for improving the generalizability of unseen drug compounds. In contrast, as our dataset does not contain any CYP information associated with food compounds, ArcDFI has limited generalizability in unseen foods because its Compound-CYP interaction block relies solely on traditional unsupervised cross-attention mechanisms.

Furthermore, our substructure-wise encoding method builds dense embeddings that generalize robustly to novel drug scaffolds. This is also demonstrated by SSI-DDI, a DDI prediction model that incorporates the concept of molecular substructures, which showed the second-highest AUROC, AUPRC, and F1 scores under the cold drug setting.

In contrast, the cold food experiments reveal DeepDrug outperforming SSI-DDI, which we attribute to the generally simpler scaffolds of many food compounds—characterized by fewer heavy atoms—where substructure-based features may lose representation power. Contrary to SSI-DDI and our ArcDFI, DeepDrug builds atom-level and bond-level embeddings derived from SMILES sequences. This allows it to better capture fine-grained atomic context and local chemical environments, which can be especially advantageous when modeling interactions involving smaller, structurally simple molecules. These observations highlight the

Table 3. Evaluation results for ArcDFI, its ablated version, and baseline models under the cold drug experiment setting. All evaluation scores were averaged over three iterations along with their standard deviation. Best results are bold-faced.

Model	Cold Drug Evaluation Setting				
	AUROC	AUPRC	F1	Precision	Recall
DeepSynergy	0.8638±0.0006	0.8122±0.0047	0.6591±0.0061	0.5155±0.0080	0.9138±0.0027
DeepDDI	0.9048±0.0350	0.9151±0.0511	0.7982±0.0398	0.7078±0.0532	0.9151±0.0196
EPGCN-DS	0.9050±0.0001	0.8454±0.0003	0.7125±0.0057	0.8411±0.0046	0.6181±0.0111
CASTER	0.7660±0.0058	0.6478±0.0083	0.5861±0.0158	0.4381±0.0232	0.8872±0.0226
SSI-DDI	0.9697±0.0058	0.9451±0.0124	0.8816±0.0155	0.8884±0.0096	0.8757±0.0370
MatchMaker	0.8883±0.0390	0.8472±0.0704	0.6685±0.0487	0.5258±0.0636	0.9237±0.0065
MR-GNN	0.9674±0.0017	0.9388±0.0022	0.8733±0.0092	0.8936±0.0139	0.8546±0.0292
DeepDrug	0.9615±0.0021	0.9364±0.0066	0.8551±0.0130	0.9053±0.0167	0.8111±0.0341
DFinder	0.5002±0.0094	0.0011±0.0001	0.0022±0.0001	0.0011±0.0001	0.9446±0.0682
DFI-MS	0.7095±0.0111	0.5614±0.0097	0.5543±0.0023	0.3917±0.0011	0.9480±0.0065
ArcDFI ($-\mathcal{L}_2$)	0.9474±0.0134	0.8762±0.0517	0.8385±0.0153	0.7775±0.0285	0.9104±0.0104
ArcDFI (ours)	0.9717±0.0034	0.9494±0.0082	0.8833±0.0056	0.8761±0.0196	0.8911±0.0165

<https://doi.org/10.1371/journal.pcbi.1013055.t003>

Table 4. Evaluation results for ArcDFI, its ablated version and baseline models under the cold food experiment setting. All evaluation scores were averaged over three iterations along with their standard deviation. Best results are bold-faced.

Model	Cold Food Evaluation Setting				
	AUROC	AUPRC	F1	Precision	Recall
DeepSynergy	0.8013±0.0047	0.6070±0.0159	0.5956±0.0055	0.4612±0.0078	0.8405±0.0040
DeepDDI	0.7624±0.0102	0.4940±0.0195	0.6074±0.0081	0.4697±0.0105	0.8596±0.0176
EPGCN-DS	0.8665±0.0021	0.7031±0.0154	0.6687±0.0054	0.6985±0.0039	0.6415±0.0131
CASTER	0.8961±0.0074	0.7974±0.0319	0.6958±0.0064	0.5798±0.0128	0.8701±0.0099
SSI-DDI	0.9311±0.0071	0.8442±0.0291	0.7938±0.0086	0.7896±0.0087	0.7986±0.0253
MatchMaker	0.7620±0.0516	0.5829±0.1023	0.5395±0.0135	0.3962±0.0056	0.8459±0.0415
MR-GNN	0.9219±0.0021	0.8046±0.0140	0.7791±0.0134	0.7859±0.0383	0.7766±0.0599
DeepDrug	0.9366±0.0038	0.8565±0.0176	0.7821±0.0107	0.8574±0.0083	0.7191±0.0169
DFinder	0.2870±0.0072	0.0003±0.0000	0.0010±0.0000	0.0005±0.0000	0.9989±0.0006
DFI-MS	0.6061±0.0280	0.3905±0.0405	0.4680±0.0050	0.3138±0.0041	0.9202±0.0095
ArcDFI ($-\mathcal{L}_2$)	0.9462±0.0069	0.8586±0.0155	0.8129±0.0094	0.7652±0.0086	0.8673±0.0207
ArcDFI (ours)	0.9321±0.0073	0.8074±0.0362	0.8029±0.0136	0.7648±0.0442	0.8471±0.0231

<https://doi.org/10.1371/journal.pcbi.1013055.t004>

strength of our shared compound encoding strategy while also suggesting that incorporating localized, atom-level features may further enhance generalizability to novel food-derived compounds.

Moreover, our model design approach, which features protein language modeling for CYP450 isoenzymes, the mixture of two compound modalities (molecular graph, substructures) and efficient parameter sharing between two compound inputs (Chemical Substructure Embedding Layer, Compound Graph Encoder, CYP Interaction Block), proved to be the optimal choice for this DFI prediction task. Whereas the other baseline models did not utilize CYP450 information or cross-modality fusion, our model architecture exhibited superiority in both cold drug and food evaluation settings.

Although DFinder and DFI-MS were expected to perform better than the DDI prediction models, they exhibited poor performance when evaluated in the new drug and food setting. The poor performances of both models are related to their inherent model structures. DFinder is a graph network-based model that relies on network topology for learning compound features represented as node embeddings. Similarly, DFI-MS depends on a set of embeddings that are only partially updated during training, leaving the model unable to generalize well to novel compounds, resulting in reduced performance in the evaluated settings. We remark that these two models require to be retrained on newly introduced compounds but ours can be directly utilized by running inference on such cases. Under this strict evaluation setting, our ArcDFI model was able to overall demonstrate robustness to the unfamiliar drug and food compounds, highlighting its potential to profile newly developed medications for food interaction risks.

Analysis on attention weights

To investigate the underlying relationships between CYP450 isoenzymes and drug-food compounds, we performed model inference and visualized the calculated attention weights extracted from ArcDFI's Compound-CYP interaction block for each compound (drug, food) and head (substrate, inhibition) using heatmaps, as shown in Fig 4a, 4b, 4e, 4f and Fig 5a, 5b, 5e, 5f. Red cells indicate the pairwise interaction scores between the drug compound and each CYP450 isoenzyme under a specific interaction type (substrate or inhibition). Note that the pairwise scores were calculated based on the summation of attention weights distributed

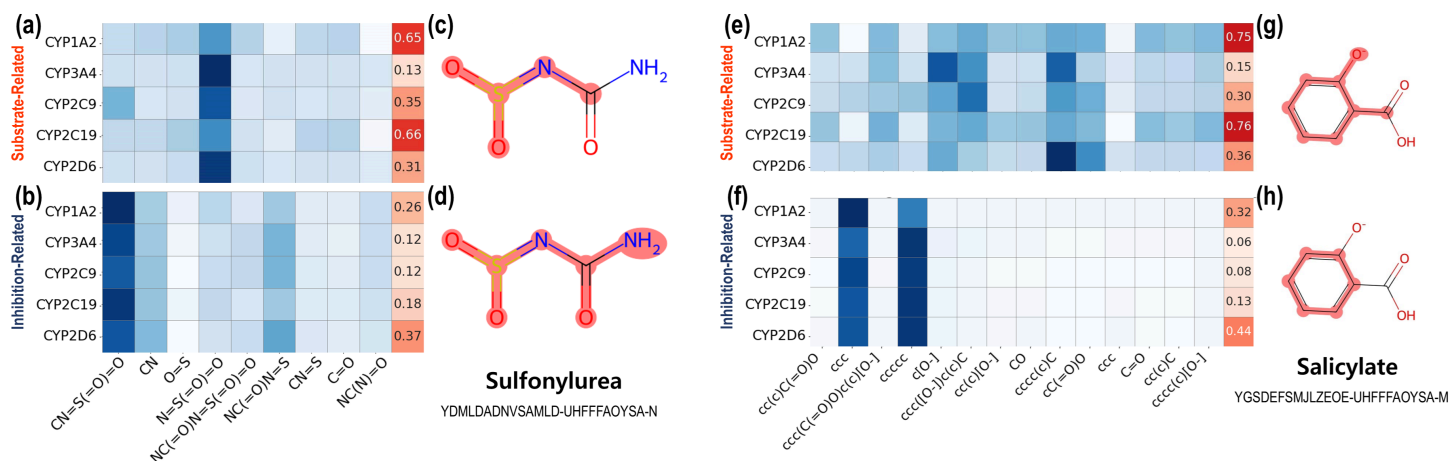


Fig 4. Analysis on ArcDFI's Compound-CYP Interaction Block for drug-food compound pair Sulfonylurea and Salicylate. (a), (b): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to substrate-type interactions for Sulfonylurea, respectively. (c), (d): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to inhibition-type interactions for Sulfonylurea, respectively. (e), (f): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to substrate-type interactions for Salicylate, respectively. (g), (h): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to inhibition-type interactions for Salicylate, respectively.

<https://doi.org/10.1371/journal.pcbi.1013055.g004>

to each compound substructure, with respect to each CYP450 isoenzyme. Blue cells indicate the contribution of each chemical substructure of a compound to CYP450 isoenzyme interactions. The chemical substructures in the heatmap were represented using SMILES.

Since the substructure-wise attention weights assigned to each CYP isoenzyme may not fully provide model interpretability, we visualized each compound with its important substructural regions being colored with red, as shown in Fig 4c, 4d, 4g, 4h and Fig 5c, 5d, 5g, 5h. We first extracted the top three attended substructures from our heat maps and mapped each atom onto the molecule. Finally, we highlighted the designated atoms with red color, which represent the substructural regions focused on by ArcDFI's Compound-CYP interaction block.

Fig 4 displays the attention weights and highlighted substructures extracted from the ArcDFI's Compound-CYP interaction block when given a pair of unseen drug compound and seen food compound, Sulfonylurea and Salicylate, as input. Sulfonylurea is a drug-like compound that has been traditionally used for treating diabetes, while Salicylate is an organic acid that is known to exhibit analgesic, antipyretic, and anti-inflammatory effects [37,38].

According to Fig 4a, the weights imply that Sulfonylurea forms a DCI with CYP1A2 and CYP2C19 [39] (0.65, 0.66) as a metabolized substrate while its substructure $N=S(=O)=O$ seemed to exhibit a relatively high influence on this interaction process. On the contrary in Fig 4b, Sulfonylurea exerted a relatively weaker inhibitory effect on CYP450 isoenzymes, with CYP2D6 being the most affected (0.37).

In Fig 4e, the weights indicate that Salicylate forms a DCI with CYP1A2 and CYP2C19 as substrates, which can potentially have stronger implications for a DFI between two compounds sharing the same isoenzymes as substrates (0.75, 0.76). In addition, the attention weights in Fig 4f indicate that Salicylate does not have inhibitory interactions with the CYP450 isoenzymes. Interestingly, the CYP2D6 isoenzyme received the highest score among the five isoenzymes (0.44), which was the same for Sulfonylurea.

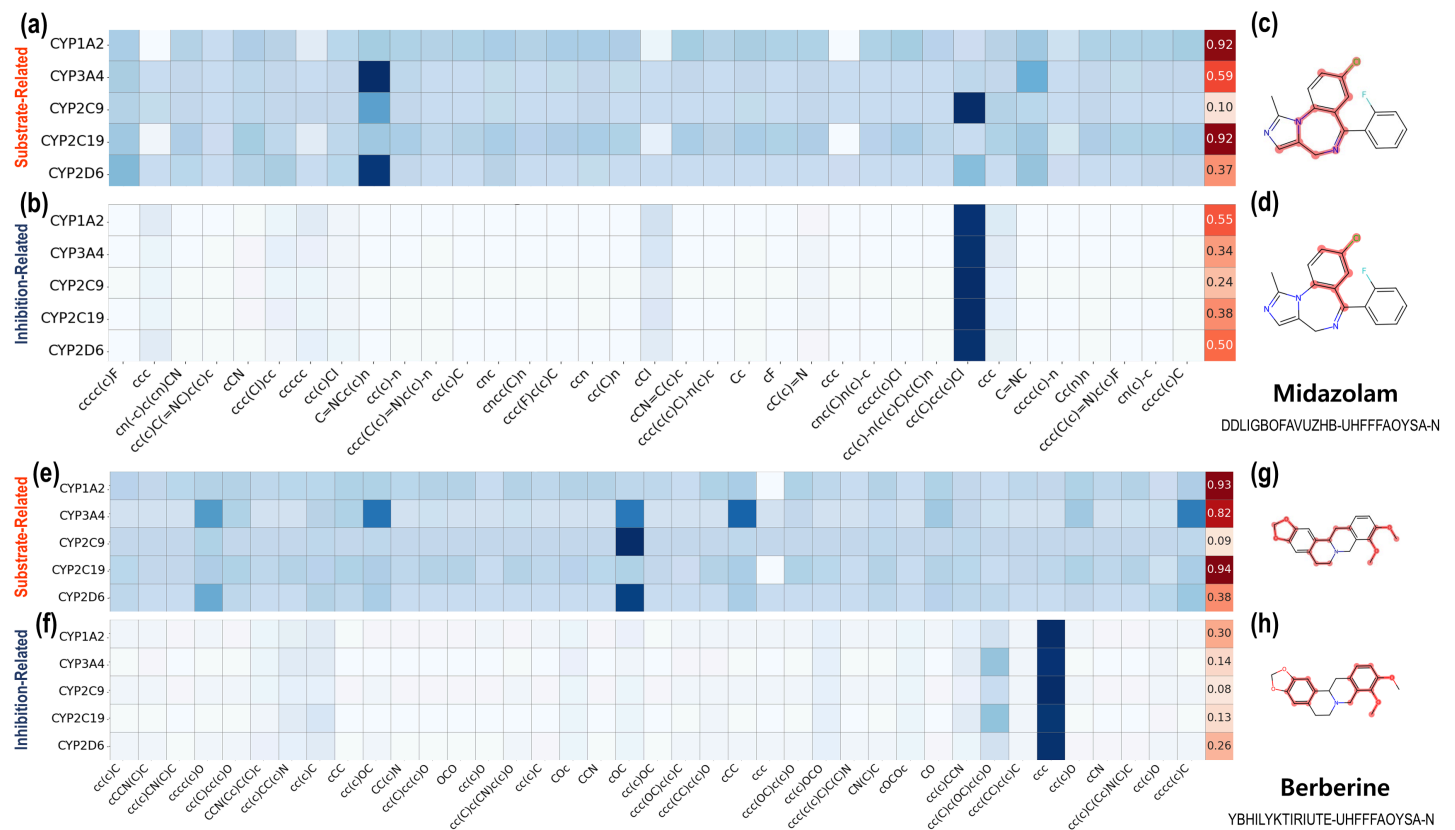


Fig 5. Analysis on ArcDFI's Compound-CYP Interaction Block for drug-food compound pair Midazolam and Berberine. (a), (b): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to substrate-type interactions for Midazolam, respectively. (c), (d): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to inhibition-type interactions for Midazolam, respectively. (e), (f): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to substrate-type interactions for Berberine, respectively. (g), (h): Attention weights and highlighted compound substructures extracted from ArcDFI's Compound-CYP Interaction Block head related to inhibition-type interactions for Berberine, respectively.

<https://doi.org/10.1371/journal.pcbi.1013055.g005>

The highlighted substructures shown in Fig 4c and 4d indicate that most of the substructural regions in Sulfonylurea are perceived to contribute to the overall attention weights, which aligns with the fact Sulfonylureas is used as a chemical derivative for making other drug compounds. As shown in Fig 4g and 4h, the aromatic substructures of Salicylate have received the most attention weights overall, which may implicate π - π interactions with the CYP450 isoenzyme targets.

Fig 5a, 5b, 5e, 5f shows the visualized attention weights when ArcDFI was given an out-of-dataset DFI pair (neither the drug nor the food compound was in our DFI dataset), with Midazolam and Berberine, as input. Midazolam is a sedative used for surgical purposes [40] and is widely used as a probe substrate for CYP3A4 isoenzyme [41], while Berberine is a plant-derived organic compound known to have anti-cancer effects [42].

As shown in Fig 5a and 5e, the attention weights indicated that both Midazolam and Berberine actively interacted with three CYP450 isoenzymes (1A2, 3A4, 2C19) as substrates. This aligns with our previous observation that ArcDFI recognizes a drug-food compound pair with an interaction when the two sides share the same CYP isoenzymes as metabolized substrates. Interestingly, no common molecular substructures exhibited similar DCI patterns, as shown in Fig 5c and 5g.

Interestingly, ArcDFI perceived Midazolam as a possible inhibitory agent targeting CYP1A2 and CYP2D6 (0.55, 0.50) as shown in Fig 5b. While this may require further investigation, we can expect ArcDFI to suggest novel DCIs and DFI predictions for unexplored drug-food pairs.

In conclusion, ArcDFI can highlight the parts of a drug or food compound most likely to interact with particular CYP450 isoenzymes. By investigating the attention weights extracted from its Compound-CYP interaction block, we can derive various DCI-related hypotheses to explain or refine the DFI predictive framework.

Discussion

Limitations

One of the main limitations of this study was the constructed DFI and DCI datasets. While the DFI dataset may seem to contain a large number of DFI pairs (609,180), it still suffers from sparsity issues as the total number of possible DFI pairs is 152,610,600, resulting in a sparsity rate of 0.4%. Despite our efforts to gather all available DCIs from multiple data sources, the number of drug compounds annotated with CYP450 isoenzyme interactions remain extremely sparse. Also, other interaction types (i.e., inducing effects) and CYP450 isoenzymes (i.e., CYP2E1) should be considered. As some studies state that Salicylate induces CYP2C19 [43], we expect our model's Compound-CYP interaction block to align with this perspective, only if its attention regularization is augmented by sufficient annotated DCI data related to induction-type interactions.

The absence of available CYP interaction labels for food compounds restricts the potential of the attention regularization method, as shown in the evaluation results for the cold food setting experiments. We suspect that, without direct supervision from annotated food-CYP interactions, the regularized attention weights introduce misleading signals that impair the food compound representations. Also, the attention weights related to the interaction-type interaction of Berberine with CYP450 isoenzymes were inconsistent with experimental studies [44,45]. This is because the Compound-CYP interaction block relies only on the DCI information for modulating its inherent attention mechanism, which lacks guidance from Food-CYP450 interaction (FCI) information. Based on our experimental results where our auxiliary loss objective helped ArcDFI achieve state-of-the-art performance in the cold-drug setting, we expect the same beneficial effects when incorporating annotated FCI information into the same auxiliary loss objective when performing cold-food experiments.

Another limitation is related to the use of compound substructures originating from the ECFPs of drug and food compounds. Although this approach is computationally efficient, the structural information represented by these substructures is not canonical owing to the inherent hashing process in ECFPs. That is, two identical embeddings may represent different molecular substructures. Additionally, the resolution of ECFPs is constrained by a radius of 2, potentially overlooking larger yet meaningful substructures. Although the compound graph encoder somewhat mitigates these potential weaknesses, there is still room for improvement in the representation of the substructures for a given compound.

The main premise of our study was the assumption that DFIs are primarily related to CYP450-mediated drug metabolism. While this assumption, implemented as an auxiliary loss objective in ArcDFI, demonstrated strong empirical performance in our cold drug experiments, many other factors influencing DFIs should be considered. Non-CYP enzymes or pathways are involved in the oxidative metabolism of food compounds [46]. Non-enzymatic mechanisms such as pH-dependent drug solubility or chelation may have associations with DFIs as well [47,48]. Certain food products can influence microbiota, which may alter the

metabolism of drug compounds and vice versa [49]. Furthermore, food can affect drug absorption by influencing the patient's gastrointestinal physiology. This necessitates annotating not only DFI types but also their underlying mechanisms, such as CYP450 isoenzymes or other relevant factors.

Another set of limitations stems from the scope of our task formulation and the practicality of ArcDFI. First, our current approach treats DFI prediction as a binary classification task—predicting the presence or absence of a drug–food interaction. While this formulation provides a simplified and tractable objective, it may not fully capture the multi-faceted nature of DFIs. An alternative could involve quantifying interaction strength or modeling interaction pathways between compounds. Moreover, one of the data sources used for training, FDMine, labels structurally dissimilar compound pairs as negative (non-interacting) examples, which may introduce excessively simplistic supervision. Ideally, negative samples should be curated based on explicit literature evidence, or the task itself could be reformulated as a drug–food interaction scoring problem rather than binary classification.

In addition, DFIs can be modulated by individual physiological factors such as gastrointestinal conditions or microbiota composition—both of which are highly dependent on a person's genetic background and lifestyle habits. This inter-individual variability poses challenges for model generalizability and may limit the model's utility in personalized drug prescription settings, where the same food–drug combination could have divergent effects across different individuals.

Finally, we conducted an external validation experiment using 29 drug–food compound pairs curated from the DDID database [50]. While constructing a fully out-of-distribution dataset proved difficult due to the comprehensive coverage of our original dataset, we ensured that all compound pairs in the external set were novel, even if the individual drugs and food compounds had appeared separately in the training data. ArcDFI achieved an accuracy of 0.7241 and an AUROC of 0.5895, indicating limited robustness. However, due to the small number of samples in the external set, the AUROC may be unstable and not fully reflective of the model's generalization ability. Nevertheless, these results still highlight the necessity of incorporating annotated food–CYP450 interaction (FCI) data to better guide the attention mechanism and improve the model's generalization not only to novel compounds but also to new compound pairings.

Future work

We propose several directions for future improvements to ArcDFI. First, we plan to utilize automated text-mining approaches leveraging Large Language Models to collect more experimentally known DFIs and CYP450 interactions from the biomedical literature, and to improve the dataset quality through manual curation performed by domain experts. Furthermore, we plan to employ molecular docking simulations and binding affinity prediction tools to augment the unlabeled DCIs in our datasets. These additional data sources enhance the utility of attention regularization, thereby improving the training and predictive performance of ArcDFI.

Second, we plan to seek ways to improve the current design of the ArcDFI model architecture, specifically related to its molecular representation learning and attention regularization method. In particular, we plan to explore better alternatives for representing compound substructures using other fingerprint-based representations, such as pharmacophores, molecular fragmentization (BRICS Decomposition [51], RECAP Algorithm [52]), or SELFIES representation [53]. Moreover, we plan to implement task-specific modifications to the attention layers

and devising robust optimization strategies to maximize the synergistic effects of attention regularization and DFI prediction.

Lastly, we plan to incorporate biological experiments to verify ArcDFI's DFI predictions and DCI hypotheses derived from its attention mechanism, when given an out-of-dataset drug-food pair as input for its model inference. This step is critical for strengthening the reliability and applicability of the model in real-world scenarios and for bridging computational predictions with clinical practice.

Conclusion

We introduce a novel DFI prediction model ArcDFI, by incorporating attention regularization guided by compound-CYP450 interactions, which offers improved generalizability and interpretability in DFI prediction. To assess our model's generalizability to unseen drug and food compounds, we conducted experiments under cold drug and food setting. The evaluation results on DFI prediction show ArcDFI's strong performance against 10 baseline approaches.

Also, the attention weights extracted from our model's Compound-CYP interaction block provided fresh insights into novel drug-food and compound-CYP450 relationships. We expect our model to facilitate discovery of novel drug-food interactions. Despite the challenges posed by dataset sparsity and the limited availability of CYP450 interaction labels, our model demonstrated promising results in terms of both predictive performance and interpretability.

Author contributions

Conceptualization: Donghyeon Park, Minji Jeon.

Data curation: Mogan Gim, Donghyeon Park.

Formal analysis: Donghyeon Park.

Funding acquisition: Mogan Gim, Jaewoo Kang, Donghyeon Park, Minji Jeon.

Investigation: Mogan Gim, Jaewoo Kang, Minji Jeon.

Methodology: Mogan Gim, Minji Jeon.

Project administration: Minji Jeon.

Resources: Jaewoo Kang, Minji Jeon.

Software: Mogan Gim, Donghyeon Park.

Supervision: Minji Jeon.

Validation: Mogan Gim, Minji Jeon.

Visualization: Mogan Gim, Donghyeon Park.

Writing – original draft: Mogan Gim, Donghyeon Park, Minji Jeon.

Writing – review & editing: Mogan Gim, Donghyeon Park, Minji Jeon.

References

1. Telessy I. Let's keep an eye on food-drug interaction. *Food Nutr J.* 2018;8:2575–7091.
2. Neuvonen PJ. Interactions with the absorption of tetracyclines. *Drugs.* 1976;11(1):45–54. <https://doi.org/10.2165/00003495-197611010-00004> PMID: 946598

3. Lilja JJ, Kivistö KT, Neuvonen PJ. Grapefruit juice-simvastatin interaction: effect on serum concentrations of simvastatin, simvastatin acid, and HMG-CoA reductase inhibitors. *Clin Pharmacol Ther.* 1998;64(5):477–83. [https://doi.org/10.1016/S0009-9236\(98\)90130-8](https://doi.org/10.1016/S0009-9236(98)90130-8) PMID: 9834039
4. Hirsh J, Fuster V, Ansell J, Halperin JL, American Heart Association, American College of Cardiology Foundation. American Heart Association/American College of Cardiology Foundation guide to warfarin therapy. *Circulation.* 2003;107(12):1692–711. <https://doi.org/10.1161/01.CIR.0000063575.17904.4E> PMID: 12668507
5. Wang T, Yang J, Xiao Y, Wang J, Wang Y, Zeng X, et al. DFinder: a novel end-to-end graph embedding-based method to identify drug-food interactions. *Bioinformatics.* 2023;39(1):btac837. <https://doi.org/10.1093/bioinformatics/btac837> PMID: 36579885
6. Wei J, Li Z, Zhuo L, Fu X, Wang M, Li K, et al. Enhancing drug-food interaction prediction with precision representations through multilevel self-supervised learning. *Comput Biol Med.* 2024;171:108104. <https://doi.org/10.1016/j.combiomed.2024.108104> PMID: 38335821
7. Guengerich FP. Cytochrome p450 and chemical toxicology. *Chem Res Toxicol.* 2008;21(1):70–83. <https://doi.org/10.1021/tx700079z> PMID: 18052394
8. Zanger UM, Schwab M. Cytochrome P450 enzymes in drug metabolism: regulation of gene expression, enzyme activities, and impact of genetic variation. *Pharmacol Ther.* 2013;138(1):103–41. <https://doi.org/10.1016/j.pharmthera.2012.12.007> PMID: 23333322
9. Bailey DG, Dresser G, Arnold JMO. Grapefruit-medication interactions: forbidden fruit or avoidable consequences?. *CMAJ.* 2013;185(4):309–16. <https://doi.org/10.1503/cmaj.120951> PMID: 23184849
10. Siwek M, Krupa AJ, Woron J. Interactions between grapefruit juice and psychotropic medications: an update of the literature and an original case series. *Expert Opin Drug Metab Toxicol.* 2024;20(5):333–45. <https://doi.org/10.1080/17425255.2024.2352468> PMID: 38721667
11. Moore LB, Goodwin B, Jones SA, Wisely GB, Serabjit-Singh CJ, Willson TM, et al. St. John's wort induces hepatic drug metabolism through activation of the pregnane X receptor. *Proc Natl Acad Sci U S A.* 2000;97(13):7500–2. <https://doi.org/10.1073/pnas.130155097> PMID: 10852961
12. Shalaby AS, Eid HH, El-Shiekh RA, Mohamed OG, Tripathi A, Al-Karmalawy AA, et al. Taming food-drug interaction risk: potential inhibitory effects of citrus juices on cytochrome liver enzymes can safeguard the liver from overdose paracetamol-induced hepatotoxicity. *ACS Omega.* 2023;8(29):26444–57. <https://doi.org/10.1021/acsomega.3c03100> PMID: 37521669
13. Knox C, Wilson M, Klinger CM, Franklin M, Oler E, Wilson A, et al. DrugBank 6.0: the DrugBank Knowledgebase for 2024. *Nucleic Acids Res.* 2024;52(D1):D1265–75. <https://doi.org/10.1093/nar/gkad976> PMID: 37953279
14. Flockhart D, Thacker D, McDonald C, Desta Z. The flockhart cytochrome P450 drug-drug interaction table. Division of Clinical Pharmacology, Indiana University School of Medicine; 2021.
15. Wei X, Zhang T, Li Y, Zhang Y, Wu F. Multi-modality cross attention network for image and sentence matching. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2020. p. 10941–50.
16. Bai B, Liang J, Zhang G, Li H, Bai K, Wang F. Why attentions may not be interpretable?. In: Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining, 2021. p. 25–34. <https://doi.org/10.1145/3447548.3467307>
17. Gim M, Choe J, Baek S, Park J, Lee C, Ju M, et al. ArkDTA: attention regularization guided by non-covalent interactions for explainable drug-target binding affinity prediction. *Bioinformatics.* 2023;39(39 Suppl 1):i448–57. <https://doi.org/10.1093/bioinformatics/btad207> PMID: 37387164
18. Lacruz-Pleguezuelos B, Piette O, Garranzo M, Pérez-Serrano D, Milešević J, Espinosa-Salinas I, et al. FooDrugs: a comprehensive food–drug interactions database with text documents and transcriptional data. *Database.* 2023;2023:baad075.
19. Rahman MM, Vadrev SM, Magana-Mora A, Levman J, Soufan O. A novel graph mining approach to predict and evaluate food-drug interactions. *Sci Rep.* 2022;12(1):1061. <https://doi.org/10.1038/s41598-022-05132-y> PMID: 35058561
20. FooDB. FooDB; 2025. <https://foodb.ca>
21. Wang N-N, Wang X-G, Xiong G-L, Yang Z-Y, Lu A-P, Chen X, et al. Machine learning to predict metabolic drug interactions related to cytochrome P450 isozymes. *J Cheminform.* 2022;14(1):23. <https://doi.org/10.1186/s13321-022-00602-x> PMID: 35428354
22. Hu W, Liu B, Gomes J, Zitnik M, Liang P, Pande V. Strategies for pre-training graph neural networks. *arXiv preprint* 2019. <https://doi.org/arXiv:1905.12265>
23. UniProt Consortium. UniProt: the Universal Protein Knowledgebase in 2023. *Nucleic Acids Res.* 2023;51(D1):D523–31. <https://doi.org/10.1093/nar/gkac1052> PMID: 36408920
24. Lin Z, Akin H, Rao R, Hie B, Zhu Z, Lu W. Language models of protein sequences at the scale of evolution enable accurate structure prediction. *bioRxiv.* 2022.

25. Lee J, Lee Y, Kim J, Kosiorek A, Choi S, Teh YW. Set transformer: a framework for attention-based permutation-invariant neural networks. In: International conference on machine learning. 2019. p. 3744–53.
26. Zhang L, Tozzo V, Higgins J, Ranganath R. Set norm and equivariant skip connections: putting the deep in deep sets. In: International Conference on Machine Learning. PMLR; 2022. p. 26559–74.
27. Preuer K, Lewis RPI, Hochreiter S, Bender A, Bulusu KC, Klambauer G. DeepSynergy: predicting anti-cancer drug synergy with Deep Learning. *Bioinformatics*. 2018;34(9):1538–46. <https://doi.org/10.1093/bioinformatics/btx806> PMID: 29253077
28. Cao X, Fan R, Zeng W. DeepDrug: a general graph-based deep learning framework for drug relation prediction. *bioRxiv*. 2020.
29. Sun M, Wang F, Elemento O, Zhou J. Structure-based drug-drug interaction detection via expressive graph convolutional networks and deep sets (student abstract). *AAAI*. 2020;34(10):13927–8. <https://doi.org/10.1609/aaai.v34i10.7236>
30. Huang K, Xiao C, Hoang T, Glass L, Sun J. CASTER: predicting drug interactions with chemical substructure representation. *AAAI*. 2020;34(01):702–9. <https://doi.org/10.1609/aaai.v34i01.5412>
31. Nyamabo AK, Yu H, Shi J-Y. SSI-DDI: substructure-substructure interactions for drug-drug interaction prediction. *Brief Bioinform*. 2021;22(6):bbab133. <https://doi.org/10.1093/bib/bbab133> PMID: 33951725
32. Kuru HI, Tastan O, Cicek AE. MatchMaker: a deep learning framework for drug synergy prediction. *IEEE/ACM Trans Comput Biol Bioinform*. 2022;19(4):2334–44. <https://doi.org/10.1109/TCBB.2021.3086702> PMID: 34086576
33. Pasa L, Navarin N, Sperduti A. Multiresolution reservoir graph neural network. *IEEE Trans Neural Netw Learn Syst*. 2022;33(6):2642–53. <https://doi.org/10.1109/TNNLS.2021.3090503> PMID: 34232893
34. Butina D. Unsupervised data base clustering based on daylight's fingerprint and tanimoto similarity: a fast and automated way to cluster small and large data sets. *J Chem Inf Comput Sci*. 1999;39(4):747–50. <https://doi.org/10.1021/ci9803381>
35. Nguyen N-Q, Jang G, Kim H, Kang J. Perceiver CPI: a nested cross-attention network for compound-protein interaction prediction. *Bioinformatics*. 2023;39(1):btac731. <https://doi.org/10.1093/bioinformatics/btac731> PMID: 36416124
36. Nguyen N-Q, Park S, Gim M, Kang J. MulinforCPI: enhancing precision of compound-protein interaction prediction through novel perspectives on multi-level information integration. *Brief Bioinform*. 2023;25(1):bbad484. <https://doi.org/10.1093/bib/bbad484> PMID: 38180829
37. Campbell RK. Glimperide: role of a new sulfonylurea in the treatment of type 2 diabetes mellitus. *Ann Pharmacother*. 1998;32(10):1044–52. <https://doi.org/10.1345/aph.17360> PMID: 9793597
38. Hirst JA, Farmer AJ, Dyar A, Lung TWC, Stevens RJ. Estimating the effect of sulfonylurea on HbA1c in diabetes: a systematic review and meta-analysis. *Diabetologia*. 2013;56(5):973–84. <https://doi.org/10.1007/s00125-013-2856-6> PMID: 23494446
39. Wang K, Yang A, Shi M, Tam CCH, Lau ESH, Fan B, et al. CYP2C19 loss-of-function polymorphisms are associated with reduced risk of sulfonylurea treatment failure in chinese patients with type 2 diabetes. *Clin Pharmacol Ther*. 2022;111(2):461–9. <https://doi.org/10.1002/cpt.2446> PMID: 34656068
40. Nordt SP, Clark RF. Midazolam: a review of therapeutic uses and toxicity. *J Emerg Med*. 1997;15(3):357–65. [https://doi.org/10.1016/s0736-4679\(97\)00022-x](https://doi.org/10.1016/s0736-4679(97)00022-x) PMID: 9258787
41. Peter J-U, Dieudonné P, Zolk O. Pharmacokinetics, pharmacodynamics, and side effects of midazolam: a review and case example. *Pharmaceuticals (Basel)*. 2024;17(4):473. <https://doi.org/10.3390/ph17040473> PMID: 38675433
42. Tillhon M, Guamán Ortiz LM, Lombardi P, Scovassi AI. Berberine: new perspectives for old remedies. *Biochem Pharmacol*. 2012;84(10):1260–7. <https://doi.org/10.1016/j.bcp.2012.07.018> PMID: 22842630
43. Chen X-P, Tan Z-R, Huang S-L, Huang Z, Ou-Yang D-S, Zhou H-H. Isozyme-specific induction of low-dose aspirin on cytochrome P450 in healthy subjects. *Clin Pharmacol Ther*. 2003;73(3):264–71. <https://doi.org/10.1067/mcp.2003.14> PMID: 12621391
44. Guo Y, Pope C, Cheng X, Zhou H, Klaassen CD. Dose-response of berberine on hepatic cytochromes P450 mRNA expression and activities in mice. *J Ethnopharmacol*. 2011;138(1):111–8. <https://doi.org/10.1016/j.jep.2011.08.058> PMID: 21920422
45. Guo Y, Chen Y, Tan Z r, Klaassen CD, Zhou H h. Repeated administration of berberine inhibits cytochromes P450 in humans. *European Journal of Clinical Pharmacology*. 2012;68:213–7.

46. Pang X, Tang C, Guo R, Chen X. Non-cytochrome P450 enzymes involved in the oxidative metabolism of xenobiotics: focus on the regulation of gene expression and enzyme activity. *Pharmacol Ther.* 2022;233:108020. <https://doi.org/10.1016/j.pharmthera.2021.108020> PMID: 34637840
47. Owens K, Argon S, Yu J, Yang X, Wu F, Lee S-C, et al. Exploring the relationship of drug bcs classification, food effect, and gastric pH-dependent drug interactions. *AAPS J.* 2021;24(1):16. <https://doi.org/10.1208/s12248-021-00667-w> PMID: 34961909
48. Bushra R, Aslam N, Khan AY. Food-drug interactions. *Oman Med J.* 2011;26(2):77–83. <https://doi.org/10.5001/omj.2011.21> PMID: 22043389
49. Farag MA, von Bergen M, Shao P, Güven EÇ. How do food and drugs interact with gut microbiota? Toward better or worse health outcomes. *J Adv Res.* 2023;52:1. <https://doi.org/10.1016/j.jare.2023.09.006> PMID: 37777321
50. Hong Y, Xu H, Liu Y, Zhu S, Tian C, Chen G, et al. DDID: a comprehensive resource for visualization and analysis of diet-drug interactions. *Brief Bioinform.* 2024;25(3):bbae212. <https://doi.org/10.1093/bib/bbae212> PMID: 38711369
51. Degen J, Wegscheid-Gerlach C, Zaliani A, Rarey M. On the art of compiling and using “drug-like” chemical fragment spaces. *ChemMedChem.* 2008;3(10):1503–7. <https://doi.org/10.1002/cmdc.200800178> PMID: 18792903
52. Lewell XQ, Judd DB, Watson SP, Hann MM. RECAP—retrosynthetic combinatorial analysis procedure: a powerful new technique for identifying privileged molecular fragments with useful applications in combinatorial chemistry. *J Chem Inf Comput Sci.* 1998;38(3):511–22. <https://doi.org/10.1021/ci970429i> PMID: 9611787
53. Krenn M, Häse F, Nigam A, Friederich P, Aspuru-Guzik A. Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Mach Learn: Sci Technol.* 2020;1(4):045024. <https://doi.org/10.1088/2632-2153/aba947>
54. Ryu JY, Kim HU, Lee SY. Deep learning improves prediction of drug-drug and drug-food interactions. *Proc Natl Acad Sci U S A.* 2018;115(18):E4304–11. <https://doi.org/10.1073/pnas.1803294115> PMID: 29666228