# Systems biology

# CSI: Contrastive data Stratification for Interaction prediction and its application to compound–protein interaction prediction

Apurva Kalia<sup>1</sup>, Dilip Krishnan<sup>2</sup>, Soha Hassoun 💿 <sup>1,3,\*</sup>

<sup>1</sup>Department of Computer Science, Tufts University, Medford, MA 02155, United States <sup>2</sup>Google Research, Cambridge, MA 02142, Unites States

<sup>3</sup>Department of Chemical and Biological Engineering, Tufts University, Medford, MA 02155, United States

\*Corresponding author. Department of Computer Science, Tufts University, Medford, MA 02155, United States. E-mail: soha.hassoun@tufts.edu

Associate Editor: Pier Luigi Martelli

#### Abstract

**Motivation:** Accurately predicting the likelihood of interaction between two objects (compound–protein sequence, user–item, author–paper, etc.) is a fundamental problem in Computer Science. Current deep-learning models rely on learning accurate representations of the interacting objects. Importantly, relationships between the interacting objects, or features of the interaction, offer an opportunity to partition the data to create multi-views of the interacting objects. The resulting congruent and non-congruent views can then be exploited via contrastive learning techniques to learn enhanced representations of the objects.

**Results:** We present a novel method, Contrastive Stratification for Interaction Prediction (CSI), to stratify (partition) a dataset in a manner that can be exploited via Contrastive Multiview Coding to learn embeddings that maximize the mutual information across congruent data views. CSI assigns a key and multiple views to each data point, where data partitions under a particular key form congruent views of the data. We showcase the effectiveness of CSI by applying it to the compound–protein sequence interaction prediction problem, a pressing problem whose solution promises to expedite drug delivery (drug–protein interaction prediction), metabolic engineering, and synthetic biology (compound–enzyme interaction prediction) applications. Comparing CSI with a baseline model that does not utilize data stratification and contrastive learning, and show gains in average precision ranging from 13.7% to 39% using compounds and sequences as keys across multiple drug–target and enzymatic datasets, and gains ranging from 16.9% to 63% using reaction features as keys across enzymatic datasets.

Availability and implementation: Code and dataset available at https://github.com/HassounLab/CSI.

## **1** Introduction

Predicting the likelihood of interaction between two objects (e.g. user-item, spectator-movie, author-paper, label-image, compound-protein, and other pairs) is a fundamental problem in Computer Science. Recommender systems, e.g. utilize methods based on matrix-factorization to predict unknown interactions between users and items (He et al. 2017, Xue et al. 2017). In network graphs, link prediction methods can anticipate potential connections between two collaborators, or authors and papers (Vamathevan et al. 2019). Image captioning is achieved by recognizing objects within an image and characterizing interactions among them (Yao et al. 2018). Predicting the interaction between a compound and protein sequence elucidates drug-protein interactions (Bagherian et al. 2021) and promiscuous enzymatic activities on substrates (Visani et al. 2021). Across various tasks, the success of interaction prediction hinges on learned representations of the interacting objects, as high-quality representations capture key features of interest. Multiple strategies have been developed in the machine-learning literature to generate compressed representations of data (Hinton et al. 2011, Bengio et al. 2013, Goodfellow et al. 2020). Importantly, the availability of multi-modal data that represent different aspects of the same object creates opportunities for multi-view learning techniques (Li *et al.* 2019), which have proven to be a powerful way to learn representations, especially in the computer vision literature (Tian *et al.* 2020a, Radford *et al.* 2021). Some such techniques attempt to minimize the distances between congruent (same object) views, while others contrast congruent and non-congruent views of the data to push away embeddings of differing data points.

When addressing the interaction prediction problem, multimodal representation learning can be applied on each object involved in the interaction. In this case, each object is embedded within its own latent space. In some tasks, deriving congruent data views is a common place task, e.g. image cropping, chrominance, and luminance for image-related tasks. However, in other cases, identifying congruent multiviews of data is challenging or non-trivial (e.g. drugs, disease, etc.). To address this issue, and to further improve on representation learning for interactions, we use stratification (data partitioning) to generate multiple views of the data and to establish congruent and non-congruent views. Contrastive learning methods can then be applied on the stratified data to enhance learning.

Received: December 15, 2022. Revised: May 10, 2023. Editorial Decision: June 3, 2023. Accepted: July 24, 2023

© The Author(s) 2023. Published by Oxford University Press.

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 reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

More specifically, we explore how the "relationship between two interacting objects" provides an opportune data stratification strategy that allows representation learning in a joint latent space. Many-to-many interaction relationships among objects allow data to be stratified into congruent views for each object-the object itself is one view and all other objects related to it are another view. For example, in a spectator-movie interaction scenario, a set of movies preferred by the spectator becomes an alternate view of the spectator. Similarly, a set of spectators could provide an alternate view on the movie. Spectator and movie embeddings can then be learned in a joint latent space. Furthermore, "features of the interaction itself" can provide alternative views on the interacting objects. For example, where and when the interaction occurs can provide information about movies and spectators. Rational stratification of the training data enables generating congruent and non-congruent views of the objects and/or their interactions. We refer to this data stratification strategy as Contrastive Stratification for Interaction Prediction, or CSI.

To demonstrate the effectiveness of CSI, we focus on the problem of compound-protein interaction prediction, a fundamental problem in biochemistry that is prominent in drug discovery (drug-protein interaction prediction) and in understanding and engineering metabolism (compound-enzyme interaction prediction). State-of-art machine-learning methods for drug-target interaction prediction have been extensively reviewed in recent survey papers (e.g. Chen et al. 2016, Zhou et al. 2019, Abbasi et al. 2021, Bagherian et al. 2021). Related deep-learning methods broadly perform two tasks: representation learning of compounds and of protein sequences, and using the learned representations to predict interactions. Molecular representations can be learned from molecular fingerprints (Feng et al. 2018, Lee et al. 2019, Lin 2020) or learned on the corresponding molecular graphs using Graph Neural Networks (GNNs) (Tsubaki et al. 2019, Nguyen et al. 2021). Deep-learning models, such as Convolutional Neural Networks (CNNs) (Lee et al. 2019), and transformers (Huang et al. 2021, Min et al. 2021) are used to generate embeddings on protein sequences. Interaction models however remain simple, where representations are concatenated, with or without attention, to predict interaction likelihood. Unlike 3D docking simulations (Decherchi and Cavalli 2020), deep-learning models allow screening a large number of putative interactions efficiently. In addition to its importance, the problem of compound-protein interaction prediction was selected to demonstrate the effectiveness of CSI because of rich available data on enzymatic interactions. Compound-enzyme interactions are derived from known biochemical reactions, and therefore information regarding the interaction itself allow us to evaluate CSI when stratifying based on interaction features.

The core idea in CSI is intuitive. Each data point is assigned a "key" and multiple views. When learning molecular representations, each "key" is the molecule itself, and the corresponding views are the molecule and a set (or subsets) of interacting sequences. Similarly, when learning sequence representations, the "key" is the sequence itself, and the corresponding views are the sequence and the set (or subsets) of interacting molecules. When stratifying by interaction feature, the "key" is the interaction feature (e.g. all reactions that perform a specific biotransformation, such as the addition of carboxyl group), and three views of each reaction (or reaction group, if the key places multiple reactions within a strata) are readily available: reactant-product pairs associated with the reaction (View 1), compound-sequence pairs (View 2), and sequences that catalyze the reaction (View 3), where the compounds are either reaction substrates or products. Other interaction features can also be selected as keys (e.g. reactions sharing homologous sequences). Views under the same key form congruent views of the data, while views across different keys become non-congruent views. Once congruent and noncongruent data views are established, it is possible to apply any contrastive learning technique to learn the joint representation. In our case, we use Contrastive Multiview Coding (CMC) (Tian et al. 2020a), which simultaneously maximizes the mutual information present among the congruent views of the data while discarding features that are not shared among the views. Importantly, our work demonstrates the importance of view selection when applying contrastive learning (Tian et al. 2020b).

We train and evaluate CSI models for three datasets. The BindingDB dataset (Gilson *et al.* 2016) contains purchasable drugs and their protein targets that exhibit an affinity higher than  $10 \,\mu$ M, and is larger and more diverse than earlier drugprotein interaction datasets. The BRENDA dataset is derived from the BRENDA database (Chang *et al.* 2021), which provides continued manual and automated curation on enzymes and compounds interacting with enzymes. The KEGG dataset is derived from the KEGG database (Kanehisa *et al.* 2021), which catalogues biochemical reactions for a large set of organisms. The contributions of this article are:

- A generalizable data stratification method, CSI, for view selection on interacting objects, where stratification is applied either on each of the items involved in the interaction in the context of the other object, or on features of the interaction itself.
- Congruent and non-congruent data views allow CSI to be paired with contrastive learning schemes, such as CMC, resulting in learned embeddings suited for downstream tasks.
- Demonstrating how CSI applies to the compound-protein interaction prediction task for protein-drug and to enzyme-compound datasets. The latter dataset is rich in auxiliary interaction information that lends itself to stratification on interaction features.
- Showing that CSI significantly outperforms a baseline model that does not use CSI, where average precision (AP) is improved by 18.2% on the BindingDB dataset, 39% on the BRENDA dataset, and 13.7% on the KEGG dataset, when stratifying by compound and by sequence. When stratifying by reaction features for the KEGG dataset, an AP improvement of 16.9% is achieved over baseline, thus outperforming stratification by compound and by sequence.

#### 2 Materials and methods

# 2.1 Stratification on interaction data—congruent views of compounds and of protein sequences

An interaction dataset consists of compound-protein pairs known to interact. A compound may interact with multiple proteins, and a protein may have interactions with multiple compounds (Fig. 1). For data stratified using compounds as



Figure 1. Many-to-many interactions between compounds and protein sequences allow data stratification by: (A) compound, and (B) sequence

the key, the set of all protein sequences that interact with the given compound presents a view congruent with the compound. Assuming a lock-and-key-based binding model (Tripathi and Bankaitis 2017), the rationale for these views being congruent is that the interacting proteins have common features that enable binding with the same compound. Subsets, or even pairs, of the protein sequences therefore offer a view that is congruent with the compound. To simplify our formulation and implementation, we use two sequences as a congruent view of a compound. Increasing that number would result in encoders with a higher number of trainable weights. Assuming that *I* is the set of known interactions on compounds *C* and a set of sequences *S*, the set of congruent views,  $V_{\rm C}$ , for all compounds in *C* is:

$$V_{C} = \{ [c, (s_{i}, s_{j})], \ s_{i}, s_{j} \in S, c \in C, \forall (c, s_{i}), (c, s_{j}) \in I \},$$
(1)

where the square brackets denote views.

Stratification using sequences as keys is used to define congruent views for each sequence. A set of compounds, or a subset thereof, presents a congruent view of a sequence. Using a pair of compounds as a congruent view of a sequence, the set of congruent views,  $V_S$ , for all sequences in S is:

$$V_{S} = \{ [s, (c_{i}, c_{j})], s \in S, c_{i}, c_{j} \in C, \forall (c_{i}, s), (c_{j}, s) \in I \}.$$
(2)

# 2.2 Stratification on reaction data—congruent views on interaction features

Compound-protein interactions within enzymatic datasets are associated with biochemical reactions. The auxiliary data available on the reactions can be used as keys for stratifying by interaction features (and not by compounds and sequences as presented in the prior section). Each reaction represents a set of reactants that undergo a biochemical transformation into a set of products. Homologous enzyme sequences (e.g. enzymes from different organisms catalyze the same reaction) and multiple enzymes performing similar function can catalyze the same reaction. A biochemical reaction, *b*, is assumed to be bidirectional and can be represented as: where *R* is the set of reactants, *P* is the set of products, and *E* the set of enzyme sequences that catalyze the reaction. A reaction can therefore be defined as,  $b = \{R, E, P\}$ , where the subscripts on *R*, *E*, and *P* are omitted for clarity. Each reaction therefore lends itself to three congruent views: a list of corresponding reactant–product pairs, a list of compound–sequence interactions, where a compound maybe a reactant or a product, and a list of catalyzing sequences. The set of congruent views, *V*<sub>B</sub>, for the set of biochemical reactions, *B*, is given by:

$$V_B = \{ [(r_i, p_j), (c_i, s_k), (s_k, s_l)],$$
  

$$\forall r_i \in R, p_j \in P, c_i \in R \cup P, s_k, s_l \in E,$$
  

$$b = \{R, E, P\}, \forall b \in B\}.$$
(4)

#### 2.3 CSI on interacting objects

CMC (Tian *et al.* 2020a) arrives at data representations by learning embeddings for each view, and a function,  $h_0$ , that discriminates a congruent pair among a set of non-congruent views based on the learned embeddings. Once the embeddings are learned via encoders, their parameters are frozen and can be used for the downstream task. We adopt a similar methodology for CSI.

CSI is trained in two phases (Fig. 2). In the first phase, Phase 1A and 1B, we learn embeddings on compound views and, independently, on sequence views. In Phase 1A, for compound as key, each of the congruent views of the compounds,  $V_c$ , consist of one compound and two sequences. We therefore train encoders to generate embeddings for these two views, ensuring that they produce same-dimension embeddings. As compounds can be represented as graphs, we utilize a GNN [Graph Convolutional Networks (GCNs)] to encode the compounds:

$$z_{\nu 1,comp} = GCN(c). \tag{5}$$

For the protein sequence, we use a 1D CNNs on the encoded FASTA (Lipman and Pearson 1985) sequence, F, normalized to a fixed length (e.g. 1000). As we need to learn the representation of two sequences at-a-time to represent the compound, we utilize a Siamese CNN network that uses the same weights. The twin CNNs are trained in tandem on two encoded input sequences and compute the final embedding for the view,  $z_{v2}$ . That is,

$$z_{\nu2,comp} = \text{CNN}(F_i) \oplus \text{CNN}(F_j), \tag{6}$$

where  $\oplus$  is the concatenation operation.

In Phase 1B, for sequence as key, congruent views of a sequence,  $V_s$ , comprise one sequence, s, and two compounds,  $c_i$ and  $c_j$ . To learn the embeddings for these views, we utilize a CNN for the encoded sequence, and a Siamese GCN network for the two compounds. That is,

$$z_{\nu 1,seq} = GCN(c_i) \oplus GCN(c_j), \tag{7}$$



**Figure 2**. CSI model when stratifying each interacting object. (A) Phase 1A for compounds as keys—compound representation,  $z_{v1}$  is generated through a GCN and sequence–sequence representation,  $z_{v2}$  is generated using a Siamese CNN. (B) Similarly, in Phase 1B for sequences as keys, compound–compound representation,  $z_{v1}$ , is generated through a Siamese GCN, while sequence representation,  $z_{v2}$ , is generated through a CNN. (C) In Phase 2, the trained encoders from Phases 1A and 1B are fixed. The representations are concatenated to train an MLP for final prediction

$$z_{\nu2,seq} = CNN(s). \tag{8}$$

Independent GCNs and CNNs are trained in Phase 1A and 1B.

The discriminator function,  $h_{\theta}$ , between embeddings for pairs of *n*-th and *m*-th objects from two views v1 and v2 is defined as in prior work (Tian *et al.* 2020a) as the cosine similarity between their embeddings modulated by a temperature parameter  $\tau$ :

$$h_{\theta}(z_{\nu 1}^{n}, z_{\nu 2}^{m}) = exp\left(\frac{z_{\nu 1}^{n} \cdot z_{\nu 2}^{m}}{||z_{\nu 1}^{m}|| \cdot ||z_{\nu 2}^{m}||} \cdot \frac{1}{\tau}\right),\tag{9}$$

where  $\tau$  is a hyper-parameter that controls the importance of non-congruent views in pushing the embeddings apart in the latent space. We define the contrastive loss over a batch of size *k* as:

$$L_{contrastive}^{V_1, V_2} = \frac{1}{k} \sum_{n=1}^{k} \left[ -\mathbb{E} \left[ \log \frac{h_{\theta}(z_{\nu 1}^n, z_{\nu 2}^n)}{\sum\limits_{m=1}^{k} h_{\theta}(z_{\nu 1}^n, z_{\nu 2}^m)} \right] \right].$$
(10)

Defining the contrastive loss in the context of a batch facilitated the CSI implementation and avoided complex strategies to select the non-congruent views (Tian *et al.* 2020a). In essence, we select non-congruent views within a batch, instead of considering all possible non-congruent views within the entire dataset. As the contrastive loss  $L_{contrastive}^{V_1,V_2}$  treats  $V_1$  as the anchor view and iterates over  $V_2$ , it is not symmetrical. We can similarly anchor  $V_2$  and iterate over  $V_1$  to arrive at  $L_{contrastive}^{V_2,V_1}$ . The total contrastive loss (Tian *et al.* 2020a), giving equal weight to both views, is then,

$$L(V_1, V_2) = L_{contrastive}^{V_1, V_2} + L_{contrastive}^{V_2, V_1}$$
(11)

Once the encoders are trained to minimize the loss, their parameters are frozen during Phase 2. The interaction predictor is an MLP neural network that utilizes the learned embeddings for the compound views, and for the sequence views. The interaction predictor is trained on known positive interactions and on negative interactions, which consists of randomly selected compound–sequence pairs. For the contrastive loss [Equation (10)], views from different keys within a batch (e.g. one compound and two sequences for the compound-based stratification) are taken as non-congruent, while for the final predictor training, randomly selected compound–sequence pairs are taken as negative data. The final predictor  $\hat{y}$  is given by,

$$\hat{y} = MLP((z_{\nu1,comp} \oplus z_{\nu2,comp}) \oplus (z_{\nu1,seq} \oplus z_{\nu2,seq})).$$
(12)

The prediction loss is the cross entropy loss between  $\hat{y}$  and the ground truth *y* weighted by the ratio of negative-to-positive labeled data.

#### 2.4 CSI on interaction features

When data are keyed by interaction features (Fig. 3), we apply contrastive loss on three data views: a set of compound-compound pairs representing substrates-products, a set of paired compound-sequences and a set of sequences. The framework of CSI can be easily adapted to maximize the mutual information across the three views, as was suggested for CMC (Tian et al. 2020a), and to perform interaction prediction on the concatenated learned embeddings. In the first phase of CSI, Siamese GCN and CNN networks are used to learn the compound-compound and sequence-sequence embeddings, and a GCN-CNN are used to learn the embeddings for the compound-sequence embeddings. The contrastive loss is calculated pairwise, over all the views, as defined previously [Equation (10)]. In the second phase, encoder parameters are fixed, and the embeddings from all the neural networks are concatenated and used to train an MLP for interaction prediction.





**Figure 3.** CSI model when stratifying by interaction feature. (A) Phase 1. Contrastive loss is applied to the three data views: compound–compound pairs, compound–sequence pairs, and sequence–sequence pairs to generate three embeddings,  $z_{v_1}$ ,  $z_{v_2}$ , and  $z_{v_3}$ . (B) Phase 2. Trained encoders from Phase 1 are used to generate representations for compounds and sequences. These representations are concatenated together to train an MLP for the final prediction

Table 1. Statistics	for the three	evaluation	datasets.
---------------------	---------------	------------	-----------

	BindingDB	BRENDA	KEGG
(A) Statistics for the interaction dataset			
Number of interactions	68 347	40 693	127 884
Unique compounds	29 149	8891	6087
Unique sequences	3120	13 330	21 367
Compound-to-se- quence ratio	9.34	0.67	0.28
(B) Stratification on compounds (reported	l in sequences per strata)		
Average strata size	2.18	7.2	21
Standard deviation	6.85	31.7	64.1
Maximum strata size	339	1518	2065
Minimum strata size	1	1	1
Average sharing among strata	0.01	0.01	0.06
Average Jaccard similarity among strata	0.004	0.001	0.001
(C) Stratification on sequences (reported i	n compounds per strata)		
Average strata size	20.4	3.5	5.9
Standard deviation	47.4	4.1	8.9
Maximum strata size	623	107	331
Minimum strata size	1	1	1
Average sharing among strata	0.14	0.04	0.06
Average Jaccard similarity among strata	0.003	0.001	0.001
(D) Number of positives per data split			
 Training	49 746	32 536	100 999
Validation	6142	4095	12 626
Test	7833	492.5	14 261
Uneren test	1,000	005	1(2)

 $^{a}$  (A) Base statistics. (B) Strata statistics when stratifying by compound. (C) Strata statistics when stratifying by sequence. (D) The number of positive examples for various data splits.

## **3 Experiments and results**

#### 3.1 Dataset details

Three datasets, Binding DB, BRENDA, and KEGG, were used to evaluate CSI (Table 1A). BindingDB, www.bindingdb.org/

bind/chemsearch/marvin/SDFdownload.jsp?download\\_file=/ bind/purchase\\_target\\_10000.tsv, provides interaction data for purchasable BindingDB compounds. The KEGG dataset, www.kegg.jp/kegg/download/, is processed to extract the interaction and reaction data. The BRENDA dataset, www. brenda-enzymes.org/download.php, was downloaded as a text file. We used our own tool, PerBRENDA (https://github. com/HassounLab/PER\_BRENDA) to process the entries and extract the interaction information for enzymes and substrates.

Binding DB has the highest number of compounds and the highest number of compounds per sequence (9.34 ratio), as expected from a drug-target interaction dataset. For the BRENDA dataset, we extracted interactions between enzymes and ligands as positive interactions. The listed inhibitor interactions were included as labeled negative interactions for the interaction predictor training (Visani et al. 2021). For the KEGG dataset, the interactions were extracted from reactions available in the KEGG database. The two enzymatic datasets, the BRENDA and KEGG datasets, have overlap as the BRENDA database covers enzymes interacting with both natural and non-natural substrates, while the KEGG database covers natural interactions found in living organisms. The KEGG database provides detailed information on the underlying biochemical reactions, which enabled stratification on interaction features. The two datasets have 757 compounds that had the same canonical SMILES. Of the 21 367 unique sequences in KEGG dataset, 10 948 are in the BRENDA dataset. The KEGG dataset has  $\sim 3 \times$  more interactions than the BRENDA dataset.

For compound-based stratification (Table 1B), we report the size of each strata. Within each strata, the number of views is the square of the number of sequences divided by two as CMC is applied to pairs of sequences within each strata. A large size therefore indicates rich views within the strata. To assess the overlap among the strata, we report the average sharing among the strata and their Jaccard similarities. This latter metric gives a sense of how varied the views are across keys while also considering the strata size. We similarly summarize these metrics for sequence-based stratification (Table 1C). The KEGG dataset has more average shared sequences across compound keys (0.06) compared to the others, while the BindingDB dataset had more average shared compounds across sequence keys (0.14). When considering the Jaccard score, BindingDB has the highest similarities per strata for both compound- and sequence-based stratification.

For the interaction prediction task (Table 1D), the training data consists of positive examples comprising protein-compound pairs that are known to interact. The negative examples are randomly selected compound-sequence pairs. The selection strategy of the negatives reflects nature as most compounds and proteins do not interact. For training, we used a negative-to-positive ratio as 5:1, taking care to appropriately weight the loss during training. We created two kinds of test sets. The "Test" set included both positive and negative examples taken from the same distribution as the training set. We also generated test sets with  $5\times$ ,  $10\times$ , and  $25\times$  the number of negatives as positives to evaluate the impact of negative-topositive ratio in test. To test the generalizability of the model, we created an "Unseen Test" set that comprised the 5% least frequent compounds and sequences in each dataset, which were held out from the training dataset. We assume a 1:1 negative-to-positive ratio for the Unseen Test.

#### 3.2 Baseline model

While our proposed data stratification strategy can be applied to any interaction model, we create a baseline model (Supplementary Section S1) that utilizes GNNs to encode the molecules, and CNNs to encode the sequences. Compounds represented in SMILES format are converted to a molecular graph using rdkit (Landrum 2013). For our baseline, we use node features as the atom type, atomic mass, valence, is atom in ring, formal charge, radical electrons, chirality, degree, number of hydrogens, and aromaticity. Bond features are the bond type, whether the bond is part of a ring, conjugacity, and one hot encoding of the stereo configuration of the bond. Compound embeddings are learned using a multi-layer GNN encoder. The network consists of GCNs (Kipf and Welling 2016) that aggregate information at each node. The GCNs are followed by a pooling layer and two fully connected layers. Our baseline (Supplementary Fig. S1) is the same as the GraphDTA model (Nguyen et al. 2021), which is reported to outperform other state-of-art models like Öztürk et al. (2018), He (2016), and Cichonska et al. (2017), and thus provides a strong baseline. Since GraphDTA is a regression model that predicts the binding affinity, we modified the final layer of GraphDTA to enable binary prediction as needed for our interaction prediction problem.

#### 3.3 Experimental setup

To evaluate the CSI model, we measure the model's performance in ranking positive examples ahead of negative examples, as well as the model's ability to rank a molecule or sequence that has the highest probability of interacting with a sequence or molecule, respectively. We used AP, mean average precision, and R-precision as the metrics, the details of which are available in Supplementary Section S2.

The CSI model was trained in two steps. In the contrastive learning step, the encoders generating  $z_{\nu 1}$  and  $z_{\nu 2}$  were trained using CMC on the congruent and non-congruent data views. The model was trained for 700 epochs. The best temperature  $\tau$  was found to be 0.07 (we tried a range of 0.05–0.08). Adam (Kingma and Ba 2014) was used as the optimizer. In the interaction prediction step, the training set was divided into training, validation, and test sets in ratio 8:1:1. In this step, the predictor model was trained for 200 epochs, with early stopping on validation loss. The optimizer used was Adam.

#### 3.4 Results on stratification by compounds and sequences

The results for the three datasets are reported for test set with a negative-to-positive ratio of 1:1 (Table 2A). The CSI model shows improved performance across all datasets and across all metrics. CSI significantly outperforms the baseline model that does not use CSI, where AP is improved by 18.2% on the BindingDB dataset, 39% on the BRENDA dataset, and 13.7% on the KEGG dataset, when stratifying by compound and by sequence. The improvement in MAP over baseline is maximum on the BindingDB dataset (23.8%) for test data sorted by sequences while it is maximum on KEGG (26.2%) for test data sorted by compounds. For the Unseen Test set, CSI also shows improved performance across all metrics, and across all datasets (Table 2B). Specifically, AP improvements are 2.6%, 18.2%, and 1.6% for BindingDB, BRENDA, and KEGG datasets, respectively. Clearly the quality of these datasets are different, and hence the performance. The richness of the strata within each dataset, measured by the variety (i.e. less sharing) across views, impacts how contrastive learning performs on each dataset.

		Overall		Compound				Sequence			
		AP	R-precision	MAP	R-precision	Map@3	Precision@1	MAP	R-precision	Map@3	Precision@1
(A) Test set											
BindingDB	GraphDTA	0.839	0.783	0.914	0.844	0.913	0.842	0.802	0.709	0.799	0.681
	CSI	0.992	0.971	0.996	0.993	0.996	0.993	0.993	0.993	0.994	0.993
BRENDA	GraphDTA	0.778	0.713	0.804	0.680	0.804	0.680	0.874	0.803	0.874	0.791
	CSI	0.991	0.970	0.995	0.990	0.996	0.993	0.978	0.993	0.978	0.975
KEGG	GraphDTA	0.852	0.770	0.755	0.627	0.739	0.610	0.757	0.857	0.755	0.701
	CSI	0.969	0.902	0.953	0.918	0.956	0.939	0.809	0.968	0.808	0.793
(B) Unseen test											
BindingDB	GraphDTA	0.975	0.918	0.995	0.991	0.995	0.991	0.920	0.879	0.911	0.879
0	CSI	1.000	0.992	1.000	0.999	1.000	0.999	0.997	0.995	0.997	0.995
BRENDA	GraphDTA	0.845	0.703	0.934	0.891	0.927	0.891	0.901	0.857	0.897	0.841
	CSI	0.999	0.985	1.000	1.000	1.000	1.000	0.982	1.000	0.982	0.982
KEGG	GraphDTA	0.873	0.771	0.779	0.682	0.753	0.676	0.730	0.963	0.730	0.718
	CSI	0.886	0.773	0.915	0.878	0.908	0.869	0.718	0.934	0.717	0.697
(C) Ablation study											
BindingDB	CSI Seq Strat	0.972	0.921	0.982	0.973	0.978	0.969	0.971	0.960	0.983	0.982
0	CSI Comp Strat	0.991	0.971	0.987	0.989	0.996	0.991	0.983	0.982	0.982	0.982
KEGG	CSI Seq Strat	0.960	0.894	0.942	0.901	0.952	0.931	0.807	0.958	0.808	0.791
	CSI Comp Strat	0.953	0.882	0.923	0.869	0.931	0.900	0.798	0.951	0.802	0.778

Table 2. Interaction prediction results for a negative data ratio of 1:1 for the baseline (GraphDTA with a binary predictor instead of a regressor) and CSI models for the BindingDB, BRENDA, and KEGG datasets.<sup>a</sup>

<sup>a</sup> AP and R-precision are reported for the entire dataset. MAP, mean R-precision, MAP@3, and R-precision@1 are reported for data sorted by compounds and by sequences. (A) Test set. (B) Unseen test. (C) Ablation study to determine the individual contributions of each stratification strategy against using both strategies together.

The CSI model uses embeddings learnt based on compound and sequence stratification. To determine which of the two stratification strategies contributes more to performance gains over the baseline model, we performed ablation studies on BindingDB (non-enzymatic) and KEGG (enzymatic) datasets (Table 2C). Independently, compound and sequence stratification each contribute significantly to CSI's performance over the baseline. For BindingDB, with sequence-based stratification alone, the AP drops from 0.992 (with both stratification) to 0.972. The drop is minimal (0.992–0.991) when using only compound-based stratification. These results indicate that for BindingDB, the compound-based stratification contributes maximally to the CSI model's performance. For KEGG, the AP drops from 0.969 when using both stratification strategies to 0.953 when using only compound-based stratification. The drop is lesser when switching to sequence-based stratification (0.969-0.960). This indicates that for KEGG, the sequencebased stratification contributes maximally to the CSI model's performance. The performance of the CSI model also scales well when the ratio of negative-to-positive examples in increased to mimic what happens in nature (Supplementary Section S3).

#### 3.5 Results on stratification by reaction features

For the KEGG dataset, three interaction features were used to produce three stratification strategies. The first strategy partitions the data based on enzymes catalyzing the same reaction (e.g. homologs). The second strategy divides the interactions by the underlying biotransformation pattern associated with the substrate–product pairs. KEGG classifies reactions based on this property, and each class is referred to as an RCLASS (Kotera *et al.* 2014). Multiple reactions can belong to the same class and result in similar biotransformations. The third strategy divides the interaction data by the Enzyme **Table 3.** AP results on stratification for the baselines (no stratification, compound/sequence stratification) and by the three interaction features: reaction, RCLASS, and EC.<sup>a</sup>

Model	Test	Test (5:1)	Unseen test
(A) Summary of prior results			
Baseline (no stratification) Compound/sequence stratification	0.852 0.969	0.587 0.906	0.873 0.886
(B) Interaction features			
Reaction $(V_1, V_2, V_3)$ RCLASS $(V_1, V_2, V_3)$ EC $(V_1, V_2, V_3)$	0.989 0.990 0.988	0.963 0.913 0.943	0.982 0.954 0.979
(C) Ablation study			
Reaction Strat $(V_1, V_2)$ Reaction Strat $(V_2, V_3)$ Reaction Strat $(V_1, V_3)$	0.980 0.962 0.983	0.874 0.751 0.902	0.941 0.904 0.947

<sup>a</sup> The three views,  $V_1$ ,  $V_2$ , and  $V_3$ , correspond to substrate–product pairs, compounds–sequences, and pairs of sequences. The ablation study considers only two of the views at a time.

Commission (EC) number associated with the interaction. EC numbers provide hierarchical classification on enzymes and are represented as four numbers separated by periods (e.g. L-lactate dehydrogenase is assigned EC number 1.1.1.27). Each such EC number is associated with one or more biochemical reactions. The three keys used to partition the KEGG interaction data are therefore: the reaction, RCLASS, and the EC numbers. Further details and analysis are provided in Supplementary Section S4.

The results are reported using AP (Table 3) as this metric was well correlated in earlier analysis with other metrics. AP was reported for the baseline model on three datasets (1:1 negative-to-positive ratio, 5:1 ratio, and the Unseen Test) as

well as for the CSI model for the same datasets. All stratification strategies yield improved results over the baseline model and stratification by compound and sequence across all test sets, where stratification by reaction outperforms the baseline by 16.9%, 62.6%, and 13% on the 1:1, 5:1, and Unseen Test sets, respectively. Comparing with stratification by compound and sequence, stratification by reaction yields AP improvements over the compound and sequence stratification by 2.1%, 6.3%, and 10.8% on the 1:1, 5:1, and Unseen Test sets, respectively.

To evaluate how each of the views contributes to improvements over the baseline, we perform an ablation study (Table 3B). As stratification by reaction resulted in higher performance over RCLASS- and EC-based stratification, the ablation study is applied to the reaction-based stratification model. The model was successively trained on each combination of two views (instead of all 3). Removing  $V_1$  (substrateproduct view) contributed the most, when compared to removing the other two views, in reducing model performance, e.g. for the 5:1 positive-to-negative Test set, the AP performance is reduced from 0.963 to 0.751. The substrate-product view therefore contributes the most to the CSI model performance when stratifying by reaction features. We conjecture that the high similarity between substrate-product pairs contributes to higher mutual information when compared to the other views.

### 4 Conclusion

CSI is a generalizable data stratification technique that exploits relationships among interacting objects to define congruent (and non-congruent) views. Paired with CMC, CSI learns representations that maximize the mutual information among congruent views, leading to enhanced representations for the downstream interaction prediction task. In addition to advancing the state-ofthe-art in interaction prediction in the broader field of deep learning, CSI also advances interaction prediction between protein and molecules, as evidenced by our results. CSI was applied to three compound-protein sequence datasets, involving both enzyme-molecule and protein-target datasets. Our results show significant improvement in AP, in the range of 13.7%-39% over comparable baselines that do not utilize stratification. We further demonstrated that, for our datasets, stratification by interaction features results in improved performance over stratification on object relationships. Data stratification as described herein is the new paradigm of "Data-Centric AI" (https://spec trum.ieee.org/andrew-ng-data-centric-ai), where data stratification methods will complement advances in deep learning. A variety of contrastive learning methods, including CSI, have the potential to further advance protein-ligand interaction predictions.

#### Supplementary data

Supplementary data are available at *Bioinformatics* online.

#### **Conflict of interest**

None declared.

## Funding

This work was supported by the National Science Foundation [Award 1909536]. This work was also supported by the NIGMS of the National Institutes of Health [Award R01GM132391, R35GM148219]; and Army Research Office, MURI program [contract W911NF2210239]. The content is solely the responsibility of the authors and does not represent the views of the NSF nor the NIH.

#### References

- Abbasi K, Razzaghi P, Poso A *et al*. Deep learning in drug target interaction prediction: current and future perspectives. *Curr Med Chem* 2021;28:2100–13.
- Bagherian M, Sabeti E, Wang K *et al.* Machine learning approaches and databases for prediction of drug–target interaction: a survey paper. *Brief Bioinform* 2021;22:247–69.
- Bengio Y, Courville A, Vincent P. Representation learning: a review and new perspectives. IEEE Trans Pattern Anal Mach Intell 2013;35: 1798–828.
- Chang A, Jeske L, Ulbrich S et al. BRENDA, the ELIXIR core data resource in 2021: new developments and updates. Nucleic Acids Res 2021;49:D498–508.
- Chen X, Yan CC, Zhang X *et al.* Drug-target interaction prediction: databases, web servers and computational models. *Brief Bioinform* 2016;17:696–712.
- Cichonska A, Ravikumar B, Parri E *et al.* Computational-experimental approach to drug-target interaction mapping: a case study on kinase inhibitors. *PLoS Comput Biol* 2017;13:e1005678.
- Decherchi S, Cavalli A. Thermodynamics and kinetics of drug-target binding by molecular simulation. *Chem Rev* 2020;**120**:12788–833.
- Feng Q, Dueva E, Cherkasov A *et al.* PADME: a deep learning-based framework for drug-target interaction prediction. arXiv, arXiv:1807.09741, 2018, preprint: not peer reviewed. https://doi. org/10.48550/arXiv.1807.09741.
- Gilson MK, Liu T, Baitaluk M et al. BindingDB in 2015: a public database for medicinal chemistry, computational chemistry and systems pharmacology. Nucleic Acids Res 2016;44:D1045–53.
- Goodfellow I, Pouget-Abadie J, Mirza M *et al.* Generative adversarial networks. *Commun ACM* 2020;63:139–44.
- He T. SimBoost: a read-across approach for drug-target interaction prediction using gradient boosting machines. Ph.D. Thesis, Applied Sciences, School of Computing Science, 2016.
- He X, Liao L, Zhang H et al. Neural collaborative filtering. In: Proceedings of the 26th International Conference on World Wide Web. 173–82. 2017.
- Hinton GE, Krizhevsky A, Wang SD. Transforming auto-encoders. In: International Conference on Artificial Neural Networks. 44–51. Springer, 2011.
- Huang K, Xiao C, Glass LM *et al.* MolTrans: molecular interaction transformer for drug-target interaction prediction. *Bioinformatics* 2021;37:830–6.
- Kanehisa M, Furumichi M, Sato Y et al. KEGG: integrating viruses and cellular organisms. Nucleic Acids Res 2021;49:D545–51.
- Kingma DP, Ba J. Adam: a method for stochastic optimization. arXiv, arXiv:1412.6980, 2014, preprint: not peer reviewed. https://doi.org/ 10.48550/arXiv.1412.6980.
- Kipf TN, Welling M. Semi-supervised classification with graph convolutional networks. arXiv, arXiv:1609.02907, 2016, preprint: not peer reviewed. https://doi.org/10.48550/arXiv.1609.02907.
- Kotera M, Goto S, Kanehisa M. Predictive genomic and metabolomic analysis for the standardization of enzyme data. *Perspect Sci* 2014;1:24–32. Landrum G. RDKit documentation. 2013.
- Lee I, Keum J, Nam H. DeepConv-DTI: prediction of drug-target interactions via deep learning with convolution on protein sequences. *PLoS Comput Biol* 2019;15:e1007129.

- Lin X. DeepGS: deep representation learning of graphs and sequences for drug-target binding affinity prediction. arXiv, arXiv:2003.13902, 2020, preprint: not peer reviewed. https://doi.org/10.48550/arXiv. 2003.13902.
- Lipman DJ, Pearson WR. Rapid and sensitive protein similarity searches. *Science* 1985;227:1435-41.
- Min S, Park S, Kim S et al. Pre-training of deep bidirectional protein sequence representations with structural information. IEEE Access 2021;9:123912–26.
- Nguyen T, Le H, Quinn TP *et al.* GraphDTA: predicting drug-target binding affinity with graph neural networks. *Bioinformatics* 2021; 37:1140–7.
- Öztürk H, Özgür A, Ozkirimli E. DeepDTA: deep drug-target binding affinity prediction. *Bioinformatics* 2018;34:i821–9.
- Radford A, Kim JW, Hallacy C et al. Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. 8748–63. PMLR, 2021.
- Tian Y, Krishnan D, Isola P. Contrastive multiview coding. In: European Conference on Computer Vision. 776–94. Springer, 2020a.

- Tian Y, Sun C, Poole B et al. What makes for good views for contrastive learning? Adv Neural Inf Process Syst 2020b;33:6827–39.
- Tripathi A, Bankaitis VA. Molecular docking: from lock and key to combination lock. J Mol Med Clin Appl 2017;2(1).
- Tsubaki M, Tomii K, Sese J. Compound–protein interaction prediction with end-to-end learning of neural networks for graphs and sequences. *Bioinformatics* 2019;35:309–18.
- Vamathevan J, Clark D, Czodrowski P et al. Applications of machine learning in drug discovery and development. Nat Rev Drug Discov 2019;18:463–77.
- Visani GM, Hughes MC, Hassoun S. Enzyme promiscuity prediction using hierarchy-informed multi-label classification. *Bioinformatics* 2021;37:2017–24.
- Xue H-J, Dai X, Zhang J *et al.* Deep matrix factorization models for recommender systems. In: *IJCAI*, Vol. 17. 3203–9. Melbourne, Australia, 2017.
- Yao T, Pan Y, Li Y et al. Exploring visual relationship for image captioning. In: Proceedings of the European Conference on Computer Vision (ECCV). 684–99. 2018.
- Zhou L, Li Z, Yang J *et al.* Revealing drug-target interactions with computational models and algorithms. *Molecules* 2019;24:1714.

9