A Comparison of Embedding Aggregation Strategies in Drug-Target Interaction Prediction

Dimitrios Iliadis[1*], Bernard De Baets[1], Tapio Pahikkala[2], Willem Waegeman[1]

[1*]Department of Data Analysis and Mathematical Modelling, Ghent University, Coupure links 653, Ghent, B-9000, Belgium.
[2]Department of Computing, University of Turku, Turku, 20500, Finland, Country.

*Corresponding author(s). E-mail(s): dimitrios.iliadis@ugent.be; Contributing authors: bernard.debaets@ugent.be; aatapa@utu.fi; willem.waegeman@ugent.be;

Abstract

The prediction of interactions between novel drugs and biological targets is a vital step in the early stage of the drug discovery pipeline. Many deep learning approaches have been proposed over the last decade, with a substantial fraction of them sharing the same underlying two-branch architecture. Their distinction is limited to the use of different types of feature representations and branches (multi-layer perceptrons, convolutional neural networks, graph neural networks and transformers). In contrast, the strategy used to combine the outputs (embeddings) of the branches has remained mostly the same. The same general architecture has also been used extensively in the area of recommender systems, where the choice of an aggregation strategy is still an open question. In this work, we investigate the effectiveness of three different embedding aggregation strategies in the area of drug-target interaction (DTI) prediction. We formally define these strategies and prove their universal approximator capabilities. We then present experiments that compare the different strategies on benchmark datasets from the area of DTI prediction, showcasing conditions under which specific strategies could be the obvious choice.

Keywords: Drug-target interaction prediction, Binding affinity prediction, Recommender systems, Deep learning

1 Introduction

Drug discovery is a challenging task that has consistently proven to be time-consuming and expensive [1]. As a result, pharmaceutical companies are turning towards computational methods capable of automating more stages of their drug discovery pipelines. One of the earliest stages involves the prediction of interactions between chemical compounds and biological targets, also known as drug-target interaction (DTI) prediction, replacing large-scale biological screening experiments with more efficient approaches. In this area, two groups of computational methods are distinguished: docking methods and machine learning approaches. Docking methods simulate the physical interaction of molecules and proteins in the 3D space accounting for the physical structure [2]. Machine learning approaches, on the other hand, predict drug-target interactions by learning from data, using databases that contain results of traditional high-throughput screening experiments [3].

Recently, deep learning methods have gained a lot of interest in DTI prediction. In these methods, explicit features for the compounds and proteins are available beforehand and passed to a deep neural network. Even though many approaches exist, we will focus on a quite popular sub-category of two-branch deep neural networks. As shown in Fig. 1, these architectures consist of two neural network components that encode the compound and protein features, respectively. The two embedding vectors generated by the branches are then aggregated to obtain the final prediction. A non-exhaustive list of recent approaches that consider an architecture of that kind can be found in Table 1.

Method	Compound representation	Compound branch	Protein representation	Protein branch
DeepConv-DTI [4]	Morgan	MLP	amino acid sequence	CNN
DeepDTA [5]	SMILES	CNN	amino acid sequence	CNN
Shin et al. [6]	SMILES	transformer	Fasta	CNN
MDeePred [7]	$ECFP_4$	MLP	multi-channel protein features	CNN
Chen et al. [8]	molecular graph	GNN	amino acid sequence	BERT
Torng et al. [9]	molecular graph	GCNN	protein graph	GCNN
SSnet [10]	fingerprint	MLP	backbone curvature, torsion	CNN
Tsubaki et al. [11]	molecular graph	GNN	amino acid sequence	CNN
Kang et al. [12]	SMILES	BERT	amino acid sequence	BERT
Nguyen et al. [13]	molecular graph	GCN, GAT, GIN, GAT-GCN	amino acid sequence	CNN
Mingjian et al. [14]	molecular graph	GNN	protein graph	GNN
DeepPurpose [15]	various encodings	MLP, CNN, CNN-RNN...	various encodings	MLP, CNN, CNN-RNN...

Table 1: two-branch methods

DTI prediction is closely related to recommender systems based on collaborative filtering, such as the well-known Netflix challenge [16]. In both cases, a standard dataset takes the form of triplets: $\{compound, protein, activity\}$ or $\{user, item, interaction\}$. These triplets can be arranged in a sparse matrix, and in its simplest form the prediction task is matrix completion. However, when explicit feature representations are used, three additional prediction settings become feasible, in addition to matrix completion [17]: making predictions for new compounds, new proteins or new protein-compound combinations. In DTI prediction, every method utilizes the explicit features that are available for the compounds and proteins. In contrast, similar methods in

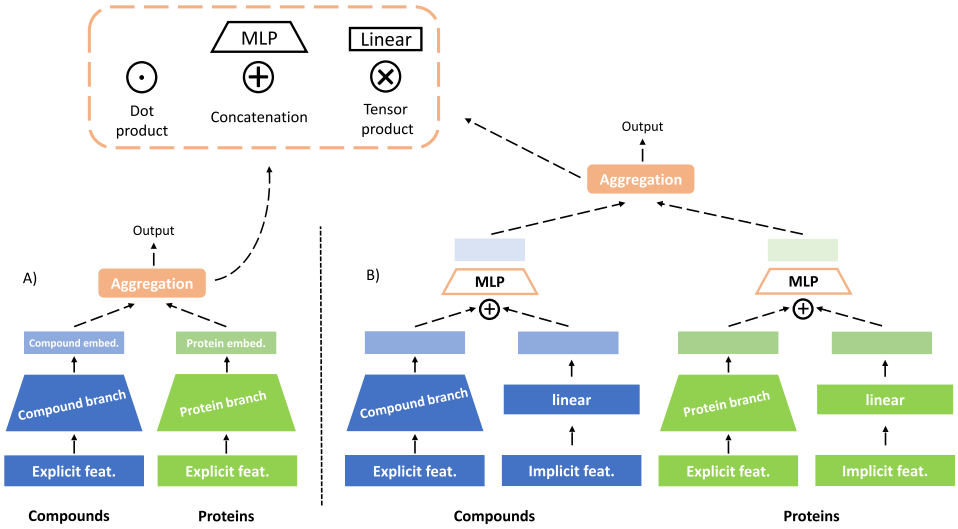

Fig. 1: A summary of the two main versions of the architecture we used. First, the two-branch architecture (sub-figure A) encodes the explicit feature representations that are available for the compounds and proteins. The resulting embeddings can be aggregated with one of the three strategies (MLP, dot product, tensor product). The second version of the architecture combines both the implicit and explicit information available for the compounds and proteins (sub-figure B). The explicit features can be encoded in the exact same manner as shown in sub-figure A, while the one-hot encoded dummy vectors for the compounds and proteins can be transformed into dense representations using a single fully-connected layer. The intermediate embeddings from the explicit and implicit features are aggregated using the MLP strategy. This output of the MLPs is the compound and protein embeddings that can then be aggregated with any of the three available strategies.

recommender systems often do not use explicit features, but create one-hot dummy vectors (implicit features) to characterize the users and items [18, 19].

DTI prediction and collaborative filtering methods also consider different aggregation strategies for the drug-target or user-item embeddings. Every DTI prediction method in Table 1 concatenates the two embeddings and then passes the resulting vector to a multi-layer perceptron (MLP strategy, see Fig. 1). Conversely, the area of collaborative filtering has a more diverse landscape, with various aggregation strategies that are regularly used, see e.g. [18, 20–26]. Especially the dot product has been extensively used in that area, and papers benchmarking the dot product versus the MLP have been published [27, 28]. Currently, a similar investigation is missing from the area of DTI prediction.

In this work, we will discuss the behavior of different embedding aggregation strategies in DTI prediction. To this end, we will analyze the above-mentioned MLP and dot product strategies, as well as a third strategy, the tensor product, which used to be

popular in the era of kernel methods [29–31]. First, we will present theoretical results highlighting the universal approximation properties of all three strategies, departing from well-known mathematical building blocks. Subsequently, we will present benchmarking results of the three strategies on combinations of DTI datasets, branch types and prediction settings. Furthermore, we will investigate the effect of adding implicit feature representations, and we will interpret the learned embeddings. As a result, the main goal of this paper is to compare the aggregation strategies in detail and not to present a winning architecture that achieves state-of-the-art performance.

2 Methods

2.1 Aggregation strategies

We first describe the details of the considered aggregation strategies. As shown in Fig. 1, we compare three strategies: the dot product, the multi-layer perceptron and the tensor product.

1. **Dot product:** In the dot product strategy, the dot product of the two embeddings is directly computed and used as the prediction of the model. This strategy requires both embeddings to have the same size, a restriction that we will further investigate in a later section.
2. **Multi-layer perceptron:** The MLP strategy concatenates the embeddings and then passes the resulting vector as input to a multi-layer perceptron, which, in turn, terminates at a single output node. As stated before, the use of an MLP increases the capacity of the overall model compared to the simple dot product operation, but, perhaps, also introduces an unwanted overhead.
3. **Tensor product:** The tensor product strategy first computes the tensor product of the two embeddings and then uses a single fully-connected layer that terminates at an output node. In fact, this aggregation strategy has never been suggested as an embedding aggregation strategy for deep neural networks, but it has been extensively used in kernel methods [29–31].

For reproducibility reasons, and to describe the universal approximation capabilities of the three aggregation strategies, we present formal definitions of the three strategies. Let $\mathcal{X}$ and $\mathcal{T}$ be two Euclidean spaces for compounds $\vec{x}$ and proteins $\vec{t}$, respectively. We formally define the problem of DTI prediction as that of estimating functions $f : \mathcal{X} \times \mathcal{T} \to \mathcal{Y}$, where $\mathcal{Y} = \mathbb{R}$ in case of regression problems. Let us consider hypothesis spaces

$$\mathcal{H}_{\mathcal{X}} = \left\{ \vec{g}_{\vec{\theta}_{\mathcal{X}}} : \mathcal{X} \to \mathbb{R}^{D_1} \mid D_1 \in \mathbb{N} \right\}$$
$$\mathcal{H}_{\mathcal{T}} = \left\{ \vec{h}_{\vec{\theta}_{\mathcal{T}}} : \mathcal{T} \to \mathbb{R}^{D_2} \mid D_2 \in \mathbb{N} \right\}$$

for learning compound embeddings from $\mathcal{X}$ to $\mathbb{R}^{D_1}$ with dimensionality D_1, and protein embeddings from $\mathcal{T}$ to $\mathbb{R}^{D_2}$ with dimensionality D_2. $\vec{\theta}_{\mathcal{X}}$ and $\vec{\theta}_{\mathcal{T}}$ denote the parameterizations of the two types of functions. Moreover, let us consider the space $C(\mathcal{X} \times \mathcal{T})$ of all continuous real-valued functions $f : \mathcal{X} \times \mathcal{T} \to \mathbb{R}$. A subspace $\mathcal{H}_{\mathrm{DP}}$ of $C(\mathcal{X} \times \mathcal{T})$

Table 2: Dataset statistics.

Dataset	#drugs	#targets	#interactions
DAVIS	68	379	30056
KIBA	2068	229	118254

corresponds to the dot product strategy of the two-branch architecture, *i.e.*, functions of the form

$$f(\vec{x}, \vec{t}) = \sum_{d=1}^{D} g_d(\vec{x})\, h_d(\vec{t})\,,$$

with $g_{\vec{\theta_\mathcal{X}}} = (g_1, \ldots, g_D) \in \mathcal{H}_\mathcal{X}$, $h_{\vec{\theta_\mathcal{T}}} = (h_1, \ldots, h_D) \in \mathcal{H}_\mathcal{T}$, and D the common dimensionality of the two embeddings.

A second subspace $\mathcal{H}_{\text{MLP}}$ corresponds to the MLP strategy of the two-branch architecture, in which the MLP is comprised of one hidden layer of size $L \in \mathbb{N}$, *i.e.*, functions of the form

$$f(\vec{x}, \vec{t}) = \vec{C}^{(3)}(\sigma \circ (\vec{B}^{(3)}[\vec{g}(\vec{x}), \vec{h}(\vec{t})] + \vec{b}^{(3)}))\,,$$

where $\vec{C}^{(3)} \in \mathbb{R}^{L \times (D_1 + D_2)}$, $\vec{B}^{(3)} \in \mathbb{R}^{1 \times L}$, $\vec{b}^{(3)} \in \mathbb{R}^{D_1 + D_2}$, and $[\vec{g}, \vec{h}]$ denotes vector concatenation of the D_1-dimensional vector $\vec{g}$ and the D_2-dimensional vector $\vec{h}$. $\sigma\circ$ represents an elementwise nonlinear transformation.

We introduce a third and final subspace $\mathcal{H}_{\text{TP}}$ that defines the tensor product strategy of the two-branch architecture, in which compound and protein embeddings are aggregated by means of a tensor product, followed by a linear layer with a single output neuron. The resulting functions are of the form

$$f(\vec{x}, \vec{t}) = \sum_{k=1}^{D_1} \sum_{l=1}^{D_2} w_{kl}\, g_k(\vec{x})\, h_l(\vec{t})\,, \tag{1}$$

where $g_{\vec{\theta_\mathcal{X}}} = (g_1, \ldots, g_{D_1}) \in \mathcal{H}_\mathcal{X}$, $h_{\vec{\theta_\mathcal{T}}} = (h_1, \ldots, h_{D_2}) \in \mathcal{H}_\mathcal{T}$, and $\vec{W} \in \mathbb{R}^{D_1 \times D_2}$.

2.2 Datasets

The proposed variants of the two-branch neural network architecture were evaluated on two benchmarks, Davis [32] and KIBA [33], because of their use as benchmarks in many of the studies presented in Table 1 and their varying numbers of compounds, proteins, and recorded affinities (see Table 2). The models were trained on the regression task, which aims to predict the affinity scores for compound-target pairs.

2.3 Prediction settings

Collaborative filtering methods usually predict missing interactions between collections of users and items that have been observed during training (random-split),

something that does not require any explicit feature representations. In contrast, the availability of such features in the typical DTI prediction task makes three additional prediction settings feasible [17]. The model could be expected to generate predictions for novel drugs (cold-drug) or novel targets (cold-target) that have not been observed during training. A fourth option that combines the strategies of the previous two is concerned with the prediction for pairs of novel drugs and targets that have not been observed during training. In the collaborative filtering area, these types of prediction settings are used less frequently but also witness a growing interest (cold start collaborative filtering). In this paper, we run experiments for the first three prediction settings: random-split, cold-drug and cold-target. The fourth prediction setting (combination of cold-drug and cold-target) was excluded from our analysis, as it is not implemented in the DeepPurpose library, which is used as the building block for protein and compound branches — see next paragraph. For each prediction setting, every dataset is split into training, validation and test sets (70-10-20%). However, the way the data is separated differs. For the cold-target setting, 70% of the targets only appear in the training dataset, 10% only appear in the validation datasets, and the remaining 20% only appear in the test set. In the cold-drug setting, the same ratios are used to split the drugs. Finally, for the random-split, the ratios are used to split the $\{compound, protein, activity\}$ triplets in the dataset.

2.4 Benchmarking experiments

For the implementation of the two-branch architectures, the DeepPurpose library [15] was chosen as the starting point. A forked version of the repository, which has been heavily modified is available online[1].

Since multiple branch pairings were possible, we included three combinations, reflecting varying degrees of descriptor and branch baseline complexity. These combinations included:

- An MLP compound branch on Morgan fingerprints [34] and an MLP protein branch on amino acid composition descriptors [35].
- A 1D Convolutional Neural Network (CNN) [36] compound branch on SMILES strings and a CNN protein branch on amino acid sequences.
- A Message-Passing Neural Network (MPNN) [37] compound branch on molecular graphs and a CNN protein branch on amino acid sequences.

2.5 Implicit feature representations

All the experiments mentioned above utilize different forms of explicit feature representations for both compounds and proteins, but not the structure of the interaction matrix. To better investigate the quality of these sources of information, we conducted additional experiments with the following differences:

- An MLP branch for the compounds and proteins that, instead of using explicit features, utilizes one-hot encoded dummy vectors. Since no generalization to new compounds or new proteins is possible when using this type of feature, we only focus

[1] https://github.com/diliadis/DeepPurpose

on the random-split setting. When the dot product is used as the aggregation strategy the resulting architecture is a close analogue to traditional matrix factorization methods.

- A two-branch architecture where each branch is comprised of an internal two-branch model (Fig. 1.B). The internal model is designed to utilize the explicit and implicit features of the compounds/proteins, something that could potentially lead to improved performance. The MLP strategy is always used for the aggregation of the internal embeddings, while all three strategies of interest are available for the external embeddings.

2.6 Hyperparameter optimization

For the comparison of the three aggregation strategies across two DTI prediction datasets and the three prediction settings, we utilized random search as the hyperparameter optimization method of choice. For every optimization round, a budget of 100 configurations was allocated, with each experiment training for up to 100 epochs (early stopping on the validation loss was also used).

The hyperparameter ranges of every experiment had to be adapted based on the aggregation strategy and branch architectures used. The full details can be found in Appendix B. Every model was trained on a single GPU (either NVIDIA Ampere A100 or NVIDIA Volta V100), and all the results were logged using the Weights and Biases platform [38].

2.7 Metrics

To guarantee a consistent comparison across all our experiments, we utilized the same regression metrics used in the majority of the work presented in Table 1. These include the mean-squared error (MSE) and r-squared (R2) metrics.

3 Results

3.1 Universal approximation

Using existing mathematical results from [39, 40] as building blocks, one can easily show that all three aggregation strategies are universal approximators. We provide further details and formal derivations in Appendix A, but summarize the main insights here. Broadly speaking, universal approximation theorems imply that neural networks can represent a wide variety of interesting functions when given appropriate weights. On the other hand, they typically do not provide a construction for the weights, but merely state that such a construction is possible.

In the setting of DTI prediction, universal approximation can only be guaranteed if the protein branch, compound branch and the aggregation strategy are universal approximators. For the formal results presented in Appendix A, we assume that the protein and compound branches are universal approximators, and we show that this is sufficient to prove universal approximation of the three aggregation strategies. For simplicity, we also assume that protein and compound feature vectors can be represented in Euclidean spaces, with multi-layer perceptrons as protein and compound

			MLP-MLP			CNN-CNN			MPNN-CNN		
Dataset	Pred setting	Aggregation	MSE	R2	Params	MSE	R2	Params	MSE	R2	Params
DAVIS	Random	Dot prod.	0.2656	0.6733	2.1	0.2622	0.6776	0.6	0.5561	0.3161	0.2
		Tensor prod.	0.2828	0.6522	6.7	0.2708	0.6669	0.4	0.5537	0.3190	0.4
		MLP	0.2722	0.6652	2.4	0.2489	0.6938	0.8	0.2647	0.6745	1.9
	Cold	Dot prod.	0.6309	-0.0426	6.3	0.6593	-0.1107	0.1	0.5112	0.1387	0.2
	Drug	Tensor prod.	0.6889	-0.1502	1.4	0.6807	-0.1467	0.3	0.5125	0.1366	0.4
		MLP	0.6763	-0.1394	4.9	0.6464	-0.0890	0.5	0.5123	0.1370	3.3
	Cold	Dot prod.	0.5895	0.3369	1.1	0.4839	0.4556	0.6	0.6811	0.2369	0.1
	Target	Tensor prod.	0.5750	0.3532	1.6	0.4803	0.4596	0.5	0.6683	0.2515	0.1
		MLP	0.5966	0.3288	2.0	0.4905	0.4482	1.7	0.5550	0.3671	0.9
KIBA	Random	Dot prod.	0.1994	0.7187	2.1	0.2169	0.6940	0.6	0.5480	0.2270	0.2
		Tensor prod.	0.2092	0.7050	4.4	0.2222	0.6865	0.4	0.5534	0.2194	0.3
		MLP	0.2093	0.7048	2.6	0.1953	0.7245	1.2	0.2544	0.6412	2.3
	Cold	Dot prod.	0.4167	0.4498	9.9	0.4730	0.3755	0.5	0.5930	0.2170	0.4
	Drug	Tensor prod.	0.4345	0.4264	7.8	0.4779	0.3690	0.4	0.5965	0.2124	0.3
		MLP	0.4258	0.4378	3.8	0.4802	0.3661	1.1	0.5172	0.3171	2.6
	Cold	Dot prod.	0.4197	0.3666	1.3	0.3839	0.4206	0.7	0.6328	0.0450	0.3
	Target	Tensor prod.	0.4126	0.3773	1.7	0.3904	0.4108	0.5	0.6326	0.0452	0.4
		MLP	0.3855	0.4181	1.8	0.3795	0.4272	0.5	0.4172	0.3703	2.6

Table 3: Performance results obtained for every possible combination of two datasets (KIBA, DAVIS), three prediction settings (random, cold-drug, cold-target), three embedding aggregation strategies and three compound-target branch pairs (MLP-MLP, CNN-CNN, MPNN-CNN). For every combination, the average MSE, R2 and number of trainable parameters of the top-5 performing (lowest overall loss) configurations are reported.

branches. Similar universal approximation theorems could be easily derived for different activation functions [41, 42], non-Euclidean spaces [43], and other types of neural network architectures, such as deep convolutional neural networks [44].

3.2 Benchmarking experiments

In this section, we present extensive comparisons of the three embedding aggregation strategies on popular benchmark datasets from the field of DTI prediction. The experiments span two different DTI prediction datasets, three prediction settings (random, cold-drug, cold-target), as well as three different combinations of input feature representations and compound-protein branch architectures. Table 3 offers a quick summary of the results that have been obtained.

In the majority of cases and for both recorded metrics, we see that the dot product and tensor product strategies can be seen as competitive alternatives to the MLP. In many cases, they achieve superior performance. At the same time, none of the three strategies can be highlighted as the strategy of choice purely based on the final performance, as none of them consistently outperforms the others. These results confirm our theoretical findings presented in Appendix A, as all three strategies can approximate any target function. Interestingly, specific combinations of dataset, prediction setting and branch pair exist, in which the three stategies result in unexpected performance differences. More specifically, when we used an MPNN as the compound branch and a CNN as the protein branch, the dot product and tensor product strategies clearly failed to reach the performance of the MLP strategy in both randomly split datasets.

A more detailed investigation of the reasons behind this result as well as a potential remedy for the dot product and tensor product strategies are presented in a later subsection.

An important characteristic of any model that can influence many practical aspects of the training process is its overall capacity. Even though capacity measures of neural networks exist (e.g. Vapnik—Chervonenkis dimension [45]), they are primarily dealing with simple MLP architectures. Since the two-branch architecture we utilize can be equipped with more complex branches (CNNs, MPNNs, etc.) we decided to simplify things and instead use the total number of trainable parameters as the measure of model capacity. Since a smaller network with fewer parameters can result in reduced memory and lower computational requirements, a smaller model that can still achieve a similar performance compared with a larger counterpart is highly desirable. Our initial hypothesis was that the overhead of the MLP strategy introduced by the fully-connected layers after the concatenation of the compound and protein embeddings would result in larger models. Based on the results shown in Table 3 the aforementioned competitiveness of the dot product strategy is usually achieved by small models. By accounting for this extra information, we can more confidently suggest the dot product strategy as a replacement for the MLP-based architecture.

Revisiting the close connection with recommender systems, a highly-cited publication by He et al. [18] first presented the dot product strategy as the simplest neural network approximation of matrix factorization. He et al. [18] then suggested the MLP strategy as a more powerful approach with the capacity to model more complex relationships between the items and users. This experimentally-backed strategy was then adopted by a series of subsequent publications [20–23] in the area of recommender systems, while proposals with the dot product strategy continued to be considered [24–26].

The superiority of the MLP aggregation strategy in the area of collaborative filtering was questioned by several subsequent publications. Rendle et al. [27] showed that, with careful hyperparameter selection, the dot product strategy could outperform the MLP strategy. They also pointed out that an MLP cannot trivially approximate the seemingly basic dot product operation. Xu et al. [28] offered a more rigorous comparison by investigating the limiting expressivity of each strategy, the convergence under the practical gradient descent algorithm, and the generalization potential. The two aforementioned publications approach the comparison of strategies exclusively in the area of recommender systems using benchmark datasets that are missing any explicit features for the users or items. In our investigation, which includes explicit feature representation and multiple generalization settings, we formulate similar conclusions as [27] and [28].

3.3 Implicit feature representations

Furthermore, Table 4 contains the results for two additional neural network configurations: two-branch neural networks that only use implicit features, and two-branch neural networks that use implicit and explicit features. So, in combination with the results of Table 3, which summarized the results for two-branch neural networks that

only included explicit features, we compare here three types of two-branch neural networks. Overall, the initial setup with only explicit features gives the best results, but the differences between the three variants is small. The negligible differences let us conclude that adding implicit features does not have benefits for the datasets and models that we considered. Hovever, the neural network that only uses implicit features still yields a satisfactory performance, so a clear structure must be present in the interaction matrices of the two datasets.

In the last decade, matrix factorization methods have been extensively used to exploit the structure in the interaction matrix by decomposing it into two small matrices that contain the implicit features [46]. Formally speaking, the structure of the interaction matrix can be summarized using the singular values of that matrix. If most singular values differ from zero, the interaction matrix has a high rank, so approximating it as a product of two smaller matrices will lead to little predictive power and meaningless implicit features. Conversely, if most singular values are equal to zero, the interaction matrix has a low rank, and low-rank matrix approximation via matrix factorization will result in a good predictive performance and meaningful inplicit features.

Matrix factorization methods share major similarities with the two-branch architectures we consider in this work. The simplest version of the two-branch architecture, which uses only the implicit features and aggregates the embeddings via the dot product strategy, can be seen as a way of performing matrix factorization [18, 28, 47–49]. For explicit feature representations and other aggregation strategies, the link with matrix factorization is less obvious, and the models become more difficult to analyze in a formal way. However, we believe that the structure of the interaction matrix is also exploited by the models in that case, because low-dimensional embeddings of proteins and compounds are constructed. To our opinion, that's the main reason why adding implicit features does not lead to performance gains in our experiments.

Let us remark that matrix factorization methods have been extensively used for DTI prediction. Early work by Cobanoglu et al. [50] used probabilistic matrix factorization combined with active learning without relying on compound or protein similarities. Ezzat et al. [51] proposed a graph regularized matrix factorization methods (and a weighted variant) to perform manifold learning and improving the performance in the cold-drug and cold-target prediction settings. More recent work by Mazzone et al. [52] used the NXTfusion [53] library that extends traditional matrix factorization methods in a non-linear fashion by inferring over an arbitrary number of data matrices, which are built as an entity relation graph. The data fusion step was performed by training a multi-task neural network that includes side information.

The concept of combining implicit and explicit information to improve the performance of DTI prediction tasks has also been incorporated into various kernel-based methods. Gonen [29] proposed a Bayesian formulation that combines kenrel-based nonlinear dimensionality reduction, matrix factorization and binary classification to predict interactions using only the chemical similarity between compounds and genomic similarity between proteins. Another kernel method, called NMTF-DTI [54], used a nonnegative tri-factorization technique based on Laplacian regularization and multiple similarity matrices for drugs and targets. A third approach defined the

Gaussian interaction profile kernel [55] to capture topological information from the drug-target interaction network, and a variation of that method improved the predictive performance even further by incorporating extra sources of chemical and genomic information with additional kernels. These results do not agree with our findings, but the datasets used in most kernel-based papers were significantly smaller than those we analyzed.

Similar to the two-branch neural networks we discuss in this work, kernel methods enable the projection of the compounds and proteins into a shared space from which the predictions are generated. However, in two-branch neural networks, the branches "learn" the compound and protein embeddings, while kernel-based methods pre-compute these embeddings by specifying a specific kernel. In general one can assume that learning the embeddings is better than defining them based on a kernel, especially when datasets are large. That's probably the reason why deep learning methods have won the battle against kernel methods in DTI prediction. Furthermore, it is worth mentioning that the use of multiple kernels for drugs and targets with the goal of improving performance [56, 57] shows many similarities with multi-modal neural network architectures that utilize multiple branches (one per feature representation) before the aggregation step [58–61]. Both concepts essentially try to boost performance by including different types of information from the same entity.

3.4 Embedding visualizations

To better understand the similarities and differences between the MLP and dot product, we experimented with various visualizations and included the most interesting examples in this work. From this point onwards, our analysis skips the tensor product, as it does not provide any performance gains or distinct characteristics compared with the MLP and dot product strategies. In their simplified form, one of the key conceptual distinctions among these strategies lies in the location of the learning process. In the dot product strategy, this is exclusively done in the two branches since the aggregation strategy is just the dot product operation. In contrast, the MLP strategy shares the learning between the two branches and the fully-connected layers after the concatenation of the embeddings. A basic visualization of the compound-protein embeddings obtained from the fully-connected layers (Fig. 2) of a well-performing MLP strategy shows an improving separation of the active and inactive compound-protein pairs the closer we get to the output node.

To investigate the quality of the compound embeddings, we experimented with another type of visualization in Figure 3. The four subplots visualize the affinity scores of four proteins (with the highest number of recorded affinities in the KIBA dataset) after the Uniform Manifold Approximation and Projection (UMAP) [62] is applied on the compound embeddings.

Even though the affinity differences observed in compound clusters between the four proteins are quite interesting, we focus on the relative comparison of the learned clusters. Based on the aforementioned discussion about the location of the learning process, the well-defined clusters of the compound embeddings from the dot product strategy are largely expected, as the two branches are exclusively responsible for learning to predict affinities. At the same time and, somewhat surprisingly, the visual

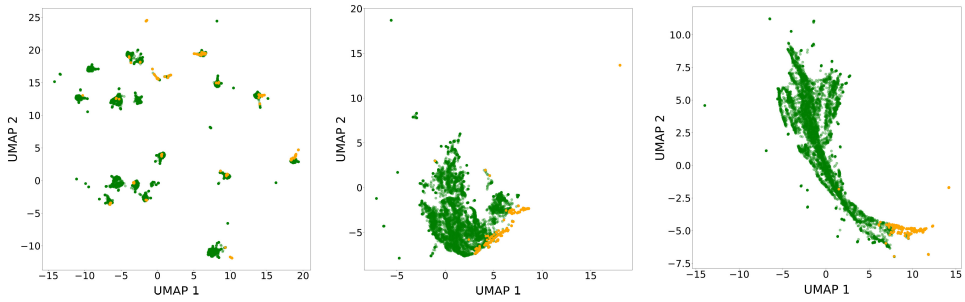

(a) First layer (b) Second layer (c) Third layer

Fig. 2: UMAP visualization of the test compound-protein embeddings extracted from the fully-connected layers of the best MLP configuration trained on the DAVIS dataset. The configuration uses CNN compound and protein branches. The points are colored using the binarized affinities (orange for active and green for inactive). From the figure, we observe that starting from the first layer (after the concatenation of the compound and protein branch embeddings) to the third, the separation of the two classes becomes clearer.

		MLP-MLP (implicit)			MLP-MLP (implicit+explicit)			CNN-CNN (implicit+explicit)		
Dataset	Aggregation	MSE	R2	Params (M)	MSE	R2	Params (M)	MSE	R2	Params (M)
DAVIS	Dot prod.	0.2804	0.6551	5.2	0.2859	0.6484	5.2	0.2690	0.6692	2.3
	Tensor prod.	0.3008	0.6301	9.6	0.3138	0.6141	9.6	0.3394	0.5825	1.3
	MLP	0.3104	0.6182	8.3	0.3217	0.6043	8.3	0.3135	0.6145	2.0
KIBA	Dot prod.	0.2481	0.6500	7.2	0.2277	0.6789	7.2	0.2296	0.6761	3.5
	Tensor prod.	0.2625	0.6297	12.9	0.2555	0.6396	12.9	0.2630	0.6290	2.6
	MLP	0.2440	0.6559	12.0	0.2590	0.6347	12.0	0.2489	0.6490	3.2

Table 4: Performance results obtained for testing combinations of implicit and explicit compound and protein feature representations. The first combination uses to MLP branches that process dummy one-hot encoded features for the compounds and proteins. The second and third combinations represent two-level two-branch architectures. The outer level is comprised of the classic two-branch architecture with embeddings that can be combined using the three strategies we investigate. The two outer branches are composed of two-branch models that utilize both explicit and implicit features. For every combination, the average MSE, R2, and number of trainable parameters of the top-5 performing (lowest overall loss) configurations are reported.

inspection of the compound embeddings obtained from the MLP strategy shows similar embedding quality as the dot product, even though the fully-connected layers after the embedding concatenation share the learning responsibility.

3.5 Oversmoothing effect and its importance when selecting an embedding aggregation strategy

For specific combinations of dataset, prediction setting and branch pair, Table 3 includes dot product and tensor product examples that are clearly outperformed by

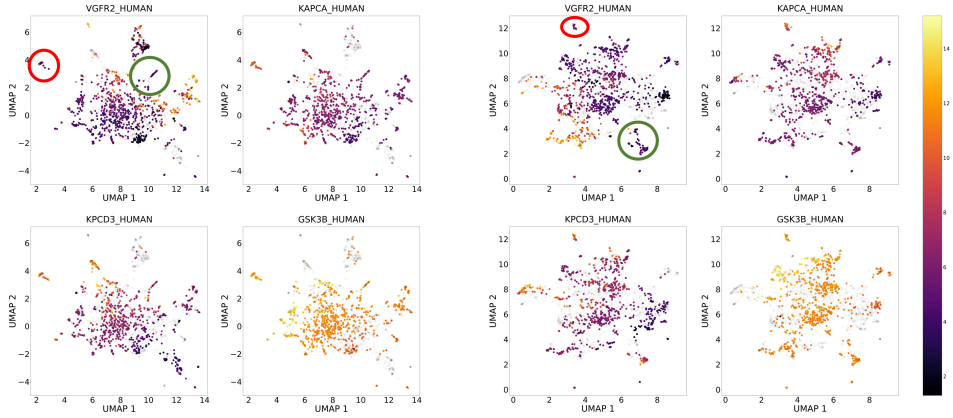

(a) MLP (b) Dot product

Fig. 3: UMAP visualization of the compound embeddings extracted from the last fully-connected layer of the compound branch (CNN) of the best MLP (left) and dot product (right) configurations trained on the KIBA dataset. The points are colored based on the affinities of the four proteins in the training set with the highest number of recorded entries. Comparing the two strategies, we observe that both create clusters of similar quality.

the MLP aggregation strategy. These cases include the use of the MPNN compound branch and the CNN protein branch on the randomly split version of the two datasets. As the CNN protein branch had been successfully used in the CNN-CNN combination, we focused our efforts on the MPNN branch as the defective model. In an effort to increase its capacity, we first tested configurations with deeper MPNNs, something that ultimately did not yield any improvements. Our hypothesis was that the over-smoothing effect that many graph neural networks suffer from [63] made any attempt to increase the capacity fail performance-wise. This is shown in the two plots of Figure 4, where points from models that use the dot product strategy and varying sizes of the MPNN branch architecture fail to approach the performance the MLP strategy achieves. Assuming that the superior performance of the MLP-based architecture was a result of the fully-connected layers, we decided to test new dot-product-based configurations by appending fully-connected layers after a smaller MPNN model and before the embedding aggregation step. With this strategy, we managed to both increase the capacity of the compound branch and avoid the over-smoothing effect that larger MPNN models suffer from. Referring back to Figure 4, these modified architectures colored in red show major improvements, as they became comparable performance-wise with the MLP strategy.

The underlying message of this exploration is that the dot product has the inability to counterbalance weaknesses that may arise from the branches, as it lacks the trainable parameters that the MLP strategy has. This problem can be tackled by increasing the capacity of the defective branch. However, cases still exist where this

strategy is inadequate because of branch-specific weaknesses (e.g., over-smoothing in GNNs). In such cases, it is suggested to increase the capacity of the defective branch with fully-connected layers, replace the branch altogether.

4 Conclusion

In this work, we analyzed alternatives to the traditional embedding aggregation strategy that has dominated the two-branch architectures in the field of DTI prediction. We presented formal and experimental results which show that all three analyzed strategies can be used to aggregate embeddings. We identified conditions under which a particular type of aggregation strategy might outperform others, and we presented various visualizations. We believe that this work can be the first step in convincing the DTI prediction community to also focus on the embedding aggregation options. Even though this may not seem vital when aggregating only two embeddings, it can become an important choice in multi-modal architectures or when more than two embeddings have to be aggregated (e.g. drug-drug-protein interaction prediction).

With regard to future work, we intend to test attention-based embedding aggregation methods, which have become popular in other application domains. Furthermore, increasing the number of embeddings that are aggregated is also an interesting avenue we intend to explore. This approach is used when different representations are available for the same entity. For example, different feature representations for a given compound (Morgan fingerprint, 2D image, molecular graph) could be combined in different ways. Thus, the order and strategy used at every stage of aggregating embeddings is a complex task. Finally, the evaluation of the added value that pre-trained embeddings from different strategies can bring to transfer learning tasks is an interesting topic that we intend to investigate in future work.

Declarations

- Conflict of interest/Competing interests: The authors declare that they have no competing interests.
- Ethics approval: Not applicable.
- Consent to participate: Not applicable.
- Consent for publication: Not applicable.
- Availability of data and materials: The KIBA and DAVIS datasets are publicly available https://github.com/futianfan/DeepPurpose_Data.
- Code availability: The scripts used to perform the experiments are also available online https://github.com/diliadis/DeepPurpose as a forked and modified repository of the original DeepPurpose project.
- Authors' contributions: Dimitrios Iliadis performed the research. All authors contributed to the manuscript, read and approved the final version.

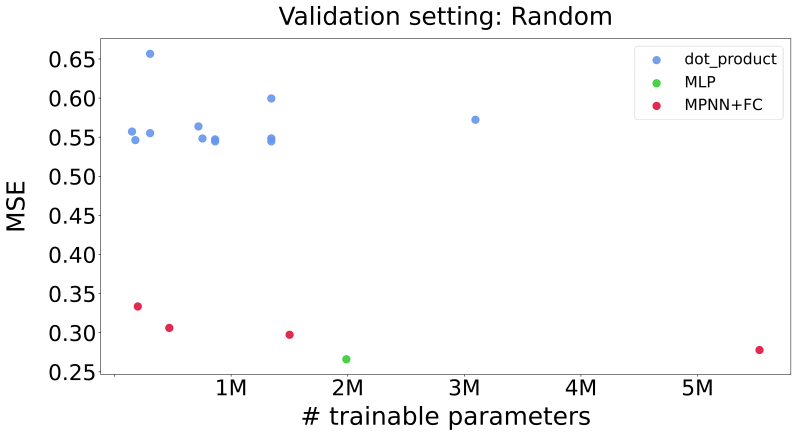

(a) random

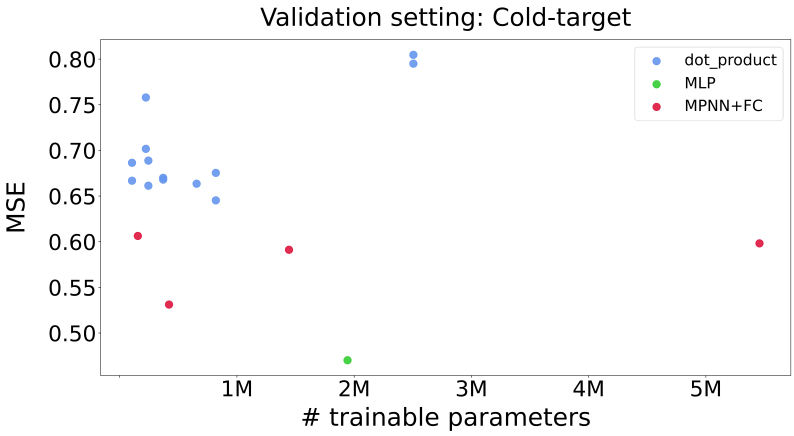

(b) cold-target

Fig. 4: Plot of the capacity of different models against their test MSE. All the points represent experiments with an MPNN as the compound branch and a CNN as the protein branch. The point colored green represents a well-performing experiment of the MLP strategy, while the blue points represent dot product strategies with similar branch configurations. The results illustrate the inability of the dot product to achieve comparable performance with the MLP strategy when one of the branches fails to utilize its input properly. Even after testing dot-product-based configurations with varying capacity (increasing the MPNN depth) exceeding that of the MLP benchmark, the performance does not improve. Our hypothesis is that the main reason for the underperforming experiments that use the dot product is the over-smoothing effect that graph neural networks suffer from. This assumption is supported by the significant improvements that a small modification attains (colored red). Instead of increasing the capacity of the MPNN directly, we instead append fully-connected layers immediately after it and before the embedding aggregation operation, thus increasing the overall capacity of the compound branch and bypassing the MPNN. The experiments colored red demonstrate that the dot product strategy can reach a comparable performance as the MLP strategy while keeping the overall capacity low.

Appendix A Universal approximation

In this appendix we formally prove that the three aggregation strategies are universal approximators, using existing mathematical results [39, 64] as building blocks. To this end, let us first formally introduce hypothesis spaces for the protein and compound branches. For simplicity, we restrict our analysis to $\mathcal{X} = \mathbb{R}^{K_1}$ and $\mathcal{T} = \mathbb{R}^{K_2}$ with $K_1, K_2 \in \mathbb{N}$, but extensions to other types of feature spaces could be easily made. For the protein branch, we will show particular interest in the hypothesis space $\mathcal{H}_{\mathcal{X},\text{MLP}}$ consisting of functions of the form

$$\vec{g}(\vec{x}) = \vec{C}^{(1)}(\sigma \circ (\vec{B}^{(1)}\vec{x} + \vec{b}^{(1)}))\,,$$

with parameterization $\vec{\theta}_{\mathcal{X},\text{MLP}} = (\vec{b}^{(1)}, \vec{B}^{(1)}, \vec{C}^{(1)}, L_1, D_1)$, where $\vec{b}^{(1)} \in \mathbb{R}^{1\times K_1}$, $\vec{B}^{(1)} \in \mathbb{R}^{L_1\times K_1}$, $\vec{C}^{(2)} \in \mathbb{R}^{D_1\times L_1}$ and $K_1, L_1, D_1 \in \mathbb{N}$. Furthermore, $\sigma\circ$ represents the sigmoid activation function.

Similarly, for the compound branch we define the hypothesis space $\mathcal{H}_{\mathcal{T},\text{MLP}}$ consisting of functions of the form

$$\vec{h}(\vec{t}) = \vec{C}^{(2)}(\sigma \circ (\vec{B}^{(2)}\vec{t} + \vec{b}^{(2)}))\,,$$

with parameterization $\vec{\theta}_{\mathcal{T},\text{MLP}} = (\vec{b}^{(2)}, \vec{B}^{(2)}, \vec{C}^{(2)}, L_2, D_2)$, where $\vec{b}^{(2)} \in \mathbb{R}^{1\times K_2}$, $\vec{B}^{(2)} \in \mathbb{R}^{L_2\times K_2}$, $\vec{C}^{(2)} \in \mathbb{R}^{D_2\times L_2}$ and $K_2, L_2, D_2 \in \mathbb{N}$.

Let us first present a theorem that establishes universal approximation for the protein and compound branches. The proof of this theorem is an immediate generalization of a classical proof given by [40] for the case $D_1 = 1$ or $D_2 = 1$.

Theorem 1. *Let $\mathcal{C}(\mathcal{X})$ and $\mathcal{C}(\mathcal{T})$ denote the sets of all continuous functions $\vec{g} : \mathcal{X} \to \mathbb{R}^{D_1}$, for any $D_1 \in \mathbb{N}$, and $\vec{h} : \mathcal{T} \to \mathbb{R}^{D_2}$, for any $D_2 \in \mathbb{N}$. Then, $\mathcal{H}_{\mathcal{X},MLP}$ and $\mathcal{H}_{\mathcal{T},MLP}$ are universal approximators of $\mathcal{C}(\mathcal{X})$ and $\mathcal{C}(\mathcal{T})$, respectively:*

(i) For all $\vec{g} \in \mathcal{C}(\mathcal{X})$ and $\epsilon > 0$ there exists $\vec{g}' \in \mathcal{H}_{\mathcal{X},MLP}$ such that

$$\sup_{\vec{x}\in\mathcal{X}} \parallel \vec{g}(\vec{x}) - \vec{g'}(\vec{x}) \parallel < \epsilon\,.$$

(ii) For all $\vec{h} \in \mathcal{C}(\mathcal{T})$ and $\epsilon > 0$ there exists $\vec{h}' \in \mathcal{H}_{\mathcal{T},MLP}$ such that

$$\sup_{\vec{t}\in\mathcal{T}} \parallel \vec{h}(\vec{t}) - \vec{h'}(\vec{t}) \parallel < \epsilon\,.$$

Universal approximation capabilities of the three types of aggregation strategies then immediately follows from existing mathematical results. To show this formally, let us first define universal approximation for DTI prediction.

Definition 1. *Let $\mathcal{C}(\mathcal{X} \times \mathcal{T})$ be the space of all continuous functions of type $f : \mathcal{X} \times \mathcal{T} \to \mathbb{R}$. A hypothesis space $\mathcal{H}$ is a universal approximator of $\mathcal{C}(\mathcal{X} \times \mathcal{T})$ if for all*

$f \in \mathcal{C}(\mathcal{X} \times \mathcal{T})$ and $\epsilon > 0$ there exists $f' \in \mathcal{H}$ such that

$$\sup_{\vec{x} \in \mathcal{X}, \vec{t} \in \mathcal{T}} |f(\vec{x}, \vec{t}) - f'(\vec{x}, \vec{t})| < \epsilon \,.$$

Let us now consider the hypothesis spaces $\mathcal{H}_{\text{MLP}}$, $\mathcal{H}_{\text{DP}}$, $\mathcal{H}_{\text{TP}}$ introduced in Section 2. We show that all three of them are universal approximators.

Theorem 2. *If $\mathcal{H}_{\mathcal{X}}$ and $\mathcal{H}_{\mathcal{T}}$ are universal approximators of $\mathcal{C}(\mathcal{X})$ and $\mathcal{C}(\mathcal{T})$, respectively, then $\mathcal{H}_{\text{DP}}$, $\mathcal{H}_{\text{MLP}}$ and $\mathcal{H}_{\text{TP}}$ are universal approximators of $\mathcal{C}(\mathcal{X} \times \mathcal{T})$.*

Proof. As discussed in [39], the universality of $\mathcal{H}_{\text{DP}}$ follows from the Stone–Weierstras theorem. The authors assume that $\mathcal{X}$ and $\mathcal{T}$ are Hausdorff spaces, but here we limit our analysis to Euclidean spaces for $\mathcal{X}$ and $\mathcal{T}$. The universality of $\mathcal{H}_{\text{TP}}$ follows from the fact that a dot product can be rewritten as a tensor product. To this end, in Eq. 1, we choose $D = D_1 = D_2$, $w_{kl} = 1$ when $k = l$ and $w_{kl} = 0$ when $k \neq l$. Similarly, the universality of $\mathcal{H}_{\text{MLP}}$ immediately follows from the fact that a dot product, like any other continuous function of several variables, can be uniformly approximated by a multi-layer perceptron. □

Appendix B Hyper-parameter ranges

In the following section, we detail the hyperparameter space we explore:

- MLP strategy:
  - layers: [1, 2, 3]
  - nodes per layer: [4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048]

and the three input branches:

- MLP compound or protein branch:
  - layers: [1, 2, 3, 4]
  - nodes per layer: [4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048]
- CNN compound or protein branch:
  - layers: [1, 2, 3, 4]
  - filters per layer: [16, 32, 64, 128]
  - kernels per layer: [4, 8, 12, 16]
- MPNN compound branch:
  - depth: [1, 2, 3]

Finally, the range for both the compound and protein embedding was the following:

$$[2, 3, 4, 8, 16, 32, 64, 128, 256, 512]$$

References

[1] Sinha, S., Vohora, D.: Drug Discovery and Development: An Overview. Pharmaceutical Medicine and Translational Clinical Research (Pharmaceut Med .), 19–32 (2018)

[2] Pujadas, G., Vaque, M., Ardevol, A., Blade, C., Salvado, M.J., Blay, M., Fernandez-Larrea, J., Arola, L.: Protein-ligand Docking: A Review of Recent Advances and Future Perspectives. Current Pharmaceutical Analysis (Curr. Pharm. Anal.) **4**(1), 1–19 (2008) https://doi.org/10.2174/157341208783497597

[3] Zanni, R., Gálvez-Llompart, M., Gálvez, J., García-Domenech, R.: QSAR multi-target in drug discovery: a review. Current Computer-Aided Drug Design (Curr. Comput) **10**(2), 129–136 (2014) https://doi.org/10.2174/1573409910021407081051 24 . Accessed 2023-05-08

[4] Lee, I., Keum, J., Nam, H.: DeepConv-DTI: Prediction of drug-target interactions via deep learning with convolution on protein sequences. PLOS Computational Biology **15**(6), 1007129 (2019)

[5] Öztürk, H., Özgür, A., Ozkirimli, E.: DeepDTA: Deep drug-target binding affinity prediction. In: Bioinformatics, vol. 34, pp. 821–829 (2018)

[6] Shin, B., Park, S., Kang, K., Ho, J.C.: Self-Attention Based Molecule Representation for Predicting Drug-Target Interaction. Proceedings of Machine Learning Research (PMLR) **106**, 1–18 (2019)

[7] Rifaioglu, A.S., Atalay, R.C., Kahraman, D.C., Doğan, T., Martin, M., Atalay, V.: MDeePred: novel multi-channel protein featurization for deep learning-based binding affinity prediction in drug discovery. Bioinformatics **37**(5), 693–704 (2021)

[8] Chen, W., Chen, G., Zhao, L., Yu-Chian Chen, C.: Predicting Drug-Target Interactions with Deep-Embedding Learning of Graphs and Sequences. The Journal of Physical Chemistry (J. Phys. Chem.) **125**, 5642 (2021)

[9] Torng, W., Altman, R.B.: Graph Convolutional Neural Networks for Predicting Drug-Target Interactions. Journal of Chemical Information and Modeling (J Chem Inf Model) (2019)

[10] Karki, N., Verma, N., Trozzi, F., Tao, P., Kraka, E., Zoltowski, B.: SSnet: A Deep Learning Approach for Protein-Ligand Interaction Prediction. International Journal of Molecular Sciences 2021 (Int. J. Mol. Sci.) **22**(3), 1392 (2021)

[11] Tsubaki, M., Tomii, K., Sese, J.: Compound–protein interaction prediction with end-to-end learning of neural networks for graphs and sequences. Bioinformatics **35**(2), 309–318 (2019)

[12] Kang, H., Goo, S., Lee, H., Chae, J.-w., Yun, H.-y., Jung, S.: Fine-tuning of bert model to accurately predict drug–target interactions. Pharmaceutics **14**(8), 1710 (2022)

[13] Nguyen, T., Le, H., Quinn, T.P., Nguyen, T., Le, T.D., Venkatesh, S.: GraphDTA: predicting drug–target binding affinity with graph neural networks. Bioinformatics **37**(8), 1140–1147 (2021)

[14] Jiang, M., Li, Z., Zhang, S., Wang, S., Wang, X., Yuan, Q., Wei, Z.: Drug–target affinity prediction using graph neural network and contact maps. RSC Advances **10**(35), 20701–20712 (2020)

[15] Huang, K., Fu, T., Glass, L.M., Zitnik, M., Xiao, C., Sun, J.: DeepPurpose: a deep learning library for drug–target interaction prediction. Bioinformatics **36**(22-23), 5545–5547 (2020)

[16] Bennett, J., Lanning, S.: The Netflix Prize. In: Proceedings of KDD Cup and Workshop (2007). https://www.semanticscholar.org/paper/The-Netflix-Prize-Bennett-Lanning/31af4b8793e93fd35e89569ccd663ae8777f0072 Accessed 2023-02-16

[17] Pahikkala, T., Airola, A., Pietilä, S., Shakyawar, S., Szwajda, A., Tang, J., Aittokallio, T.: Toward more realistic drug–target interaction predictions. Briefings in bioinformatics (Brief. Bioinformatics) **16**(2), 325–337 (2015)

[18] He, X., Liao, L., Zhang, H., Nie, L., Hu, X., Chua, T.S.: Neural collaborative filtering. In: Proceedings of the 26th International Conference on World Wide Web (WWW), pp. 173–182 (2017)

[19] Wu, Y., DuBois, C., Zheng, A.X., Ester, M.: Collaborative denoising auto-encoders for top-n recommender systems. In: Proceedings of the Ninth ACM International Conference on Web Search and Data Mining (WSDM), pp. 153–162 (2016)

[20] Chen, W., Cai, F., Chen, H., Rijke, M.D., Chen, H.: Joint Neural Collaborative Filtering for Recommender Systems. ACM Transactions on Information Systems (TOIS) **37**(4), 39 (2019)

[21] Karolina Dziugaite, G., Roy, D.M.: NEURAL NETWORK MATRIX FACTORIZATION. arXiv preprint arXiv:1511.06443 (2015)

[22] Liu, Y., Wang, S., Khan, M.S., He, J.: A novel deep hybrid recommender system based on auto-encoder with neural collaborative filtering. Big Data Mining and Analytics **1**(3), 211–221 (2018)

[23] Nguyen, D.M., Tsiligianni, E., Deligiannis, N.: Extendable Neural Matrix Completion. In: 2018 IEEE International Conference on Acoustics, Speech and Signal

Processing (ICASSP), pp. 6328–6332 (2018)

[24] Wang, T., Brovman, Y.M., Madhvanath, S.: Personalized Embedding-based e-Commerce Recommendations at eBay. arXiv preprint arXiv:2102.06156 (2021)

[25] Yang, J., Yi, X., Zhiyuan Cheng, D., Hong, L., Li, Y., Xiaoming Wang, S., Xu, T., Chi, E.H.: Mixed negative sampling for learning two-tower neural networks in recommendations. In: Companion Proceedings of the Web Conference 2020 (TheWebConf), pp. 441–447 (2020)

[26] Yi, X., Yang, J., Hong, L., Cheng, D.Z., Heldt, L., Kumthekar, A., Zhao, Z., Wei, L., Chi, E.: Sampling-bias-corrected neural modeling for large corpus item recommendations. In: Proceedings of the 13th ACM Conference on Recommender Systems (RecSys), pp. 269–277 (2019)

[27] Rendle, S., Krichene, W., Zhang, L., Anderson, J.: Neural collaborative filtering vs. matrix factorization revisited. In: Proceedings of the 14th ACM Conference on Recommender Systems (RecSys), pp. 240–248 (2020)

[28] Xu, D., Ruan, C., Korpeoglu, E., Kumar, S., Achan, K.: Rethinking Neural vs. Matrix-Factorization Collaborative Filtering: the Theoretical Perspectives. In: International Conference on Machine Learning (ICML), pp. 11514–11524 (2021). PMLR

[29] Gönen, M.: Predicting drug–target interactions from chemical and genomic kernels using bayesian matrix factorization. Bioinformatics **28**(18), 2304–2310 (2012)

[30] Stock, M., Pahikkala, T., Airola, A., De Baets, B., Waegeman, W.: A comparative study of pairwise learning methods based on kernel ridge regression. Neural computation (Neural Comput.) **30**(8), 2245–2283 (2018)

[31] Vert, J.-P., Qiu, J., Noble, W.S.: A new pairwise kernel for biological network inference with support vector machines. In: BMC Bioinformatics, vol. 8, pp. 1–10 (2007). Springer

[32] Davis, M.I., Hunt, J.P., Herrgard, S., Ciceri, P., Wodicka, L.M., Pallares, G., Hocker, M., Treiber, D.K., Zarrinkar, P.P.: Comprehensive analysis of kinase inhibitor selectivity. Nature Biotechnology (Nat. Biotechnol.) **29**(11), 1046–1051 (2011)

[33] Tang, J., Szwajda, A., Shakyawar, S., Xu, T., Hintsanen, P., Wennerberg, K., Aittokallio, T.: Making sense of large-scale kinase inhibitor bioactivity data sets: a comparative and integrative analysis. Journal of Chemical Information and Modeling (J Chem Inf Model) **54**(3), 735–743 (2014)

[34] Rogers, D., Hahn, M.: Extended-connectivity fingerprints. Journal of Chemical

Information and Modeling (J Chem Inf Model) **50**(5), 742–754 (2010)

[35] Reczko, M., Bohr, H.: The def data base of sequence based protein fold class predictions. Nucleic Acids Research (NAR) **22**(17), 3616 (1994)

[36] Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. Communications of the ACM (Commun. ACM) **60**(6), 84–90 (2017)

[37] Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: International Conference on Machine Learning (ICML), pp. 1263–1272 (2017). Proceedings of Machine Learning Research (PMLR)

[38] Biewald, L.: Experiment Tracking with Weights and Biases. Software available from wandb.com (2020). https://www.wandb.com/

[39] Allenby, P.D., Labuschagne, C.C.A.: On the uniform density of c(x) $\otimes$ c(y) in c(x $\times$ y). Indagationes Mathematicae **20**(1), 19–22 (2009) https://doi.org/10.1016/S0019-3577(09)00015-9

[40] Cybenko, G.: Approximation by superpositions of a sigmoidal function. Mathematics of Control, Signals and Systems **2**(4), 303–314 (1989)

[41] Leshno, M., Lin, V.Y., Pinkus, A., Schocken, S.: Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. Neural Networks **6**(6), 861–867 (1993) https://doi.org/10.1016/S0893-6080(05)80131-5

[42] Pinkus, A.: Approximation theory of the mlp model in neural networks. Acta Numerica **8**, 143–195 (1999)

[43] Brüel Gabrielsson, R.: Universal function approximation on graphs. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) Advances in Neural Information Processing Systems (NIPS), vol. 33, pp. 19762–19772. Curran Associates, Inc., ??? (2020). https://proceedings.neurips.cc/paper_files/paper/2020/file/e4acb4c86de9d2d9a41364f93951028d-Paper.pdf

[44] Zhou, D.-X.: Universality of deep convolutional neural networks. Applied and Computational Harmonic Analysis (ACHA) **48**(2), 787–794 (2020)

[45] Vapnik, V.N., Chervonenkis, A.Y.: On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. In: Measures of Complexity: Festschrift for Alexey Chervonenkis, pp. 11–30. Springer, Cham (2015)

[46] Waegeman, W., Dembczyński, K., Hüllermeier, E.: Multi-target prediction: a unifying view on problems and methods. Data Mining and Knowledge Discovery (KDD) **33**(2), 293–324 (2019)

[47] Chen, X., Zhang, Y., Ai, Q., Xu, H., Yan, J., Qin, Z.: Personalized key frame recommendation. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 315–324 (2017)

[48] Wang, X., He, X., Nie, L., Chua, T.-S.: Item silk road: Recommending items from information domains to social users. In: Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 185–194 (2017)

[49] Zhang, S., Yao, L., Sun, A., Tay, Y.: Deep learning based recommender system: A survey and new perspectives. ACM Computing Surveys (CSUR) **52**(1), 1–38 (2019)

[50] Cobanoglu, M.C., Liu, C., Hu, F., Oltvai, Z.N., Bahar, I.: Predicting drug–target interactions using probabilistic matrix factorization. Journal of chemical information and modeling **53**(12), 3399–3409 (2013)

[51] Ezzat, A., Zhao, P., Wu, M., Li, X.-L., Kwoh, C.-K.: Drug-target interaction prediction with graph regularized matrix factorization. IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB) **14**(3), 646–656 (2016)

[52] Mazzone, E., Moreau, Y., Fariselli, P., Raimondi, D.: Nonlinear data fusion over entity–relation graphs for drug–target interaction prediction. Bioinformatics, 348 (2023)

[53] Raimondi, D., Simm, J., Arany, A., Moreau, Y.: A novel method for data fusion over entity-relation graphs and its application to protein–protein interaction prediction. Bioinformatics **37**(16), 2275–2281 (2021)

[54] Jamali, A.A., Kusalik, A., Wu, F.: Nmtf-dti: a nonnegative matrix tri-factorization approach with multiple kernel fusion for drug-target interaction prediction. IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB) (2021)

[55] Van Laarhoven, T., Nabuurs, S.B., Marchiori, E.: Gaussian interaction profile kernels for predicting drug–target interaction. Bioinformatics **27**(21), 3036–3043 (2011)

[56] Nascimento, A.C., Prudêncio, R.B., Costa, I.G.: A multiple kernel learning algorithm for drug-target interaction prediction. BMC Bioinformatics **17**, 1–16 (2016)

[57] Zheng, X., Ding, H., Mamitsuka, H., Zhu, S.: Collaborative matrix factorization with multiple similarities for predicting drug-target interactions. In: Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD), pp. 1025–1033 (2013)

[58] Boezer, M., Tavakol, M., Sajadi, Z.: Fastdti: Drug-target interaction prediction using multimodality and transformers. In: Proceedings of the Northern Lights Deep Learning Workshop, vol. 4 (2023)

[59] Ren, Z.-H., You, Z.-H., Zou, Q., Yu, C.-Q., Ma, Y.-F., Guan, Y.-J., You, H.-R., Wang, X.-F., Pan, J.: Deepmpf: deep learning framework for predicting drug–target interactions based on multi-modal representation with meta-path semantic analysis. Journal of Translational Medicine (J. Transl. Med.) **21**(1), 1–18 (2023)

[60] Yang, X., Niu, Z., Liu, Y., Song, B., Lu, W., Zeng, L., Zeng, X.: Modality-dta: Multimodality fusion strategy for drug–target affinity prediction. IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB) **20**(2), 1200–1210 (2022)

[61] Zhou, D., Xu, Z., Li, W., Xie, X., Peng, S.: Multidti: drug–target interaction prediction based on multi-modal representation learning to bridge the gap between new chemical entities and known heterogeneous network. Bioinformatics **37**(23), 4485–4492 (2021)

[62] McInnes, L., Healy, J., Melville, J.: UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. arXiv (2020). https://doi.org/10.48550/arXiv.1802.03426 . http://arxiv.org/abs/1802.03426

[63] Rusch, T.K., Bronstein, M.M., Mishra, S.: A survey on oversmoothing in graph neural networks. arXiv preprint arXiv:2303.10993 (2023)

[64] Hornik, K., Stinchcombe, M., White, H.: Multilayer feedforward networks are universal approximators. Neural Networks **2**(5), 359–366 (1989)