ProteinPrompt: a webserver for predicting protein-protein interactions

Sebastian Canzler^{1,2,§,*}, Markus Fischer³, David Ulbricht³, Nikola Ristic³,
Peter W. Hildebrand^{3,4,5}, René Staritzbichler^{1,3,§,*}

¹Immuthera GmbH, 04275 Leipzig, Germany

²Young Investigators Group Bioinformatics and Transcriptomics,
Department of Molecular Systems Biology, Helmholtz-Centre for Environmental Research - UFZ,
Permoser Straße 15, Leipzig, 04317, Germany

³Institute of Medical Physics and Biophysics, University of Leipzig, 04107 Leipzig, Germany

⁴Charité - Universitätsmedizin Berlin, corporate member of Freie Universität Berlin and Humboldt-Universität zu Berlin,
Institute of Medical Physics and Biophysics, Charitéplatz 1, 10117 Berlin, Germany and

⁵Berlin Institute of Health at Charité - Universitätsmedizin Berlin, Charitéplatz 1, 10117 Berlin, Germany

*To whom correspondence should be addressed; E-mail: rene.staritzbichler@medizin.uni-leipzig.de

Abstract

Motivation: Protein-protein interactions play an essential role in a great variety of cellular processes and are therefore of significant interest for the design of new therapeutic compounds as well as the identification of side-effects due to unexpected binding. Here, we present **ProteinPrompt**, a webserver that uses machine-learning algorithms to calculate specific, currently unknown protein-protein interactions. Our tool is designed to quickly and reliably predict contacts based on an input sequence in order to scan large sequence libraries for potential binding partners, with the goal to accelerate and assure the quality of the laborious process of drug target identification.

Methods: We collected and thoroughly filtered a comprehensive database of known contacts from several sources, which is available as download. **ProteinPrompt** provides two complementary search methods of similar accuracy for comparison and consensus building. The default method is a random forest algorithm that uses the auto-correlations of seven amino acid scales. Alternatively, a graph neural network implementation can be selected. Additionally, a consensus prediction is available. For each query sequence, potential binding partners are identified from a protein sequence database. The proteome of several organisms are available and can be searched for contacts.

Results: To evaluate the predictive power of the algorithms, we prepared a test dataset that was rigorously filtered for redundancy. No sequence pairs similar to the ones used for training were included in this dataset. With this challenging dataset, the random forest method achieved an accuracy rate of 0.88 and an area under curve of 0.95. The graph neural network achieved an accuracy rate of 0.86 using the same dataset. Since the underlying learning approaches are unrelated, comparing the results of random forest and graph neural networks reduces the likelihood of errors. The consensus reached an accuracy of 0.89. **ProteinPrompt** is available online at:

<http://proteininformatics.org/ProteinPrompt>

The server makes it possible to scan the human proteome for potential binding partners of an input sequence within minutes. For local offline usage, we furthermore created a **ProteinPrompt** Docker image which allows for batch submission: <https://gitlab.hzdr.de/proteinprompt/ProteinPrompt>. In conclusion, we offer a fast, accurate, easy-to-use online service for predicting binding partners from an input sequence.

1 Introduction

Protein interactions are key to the complex molecular interplays of cellular processes. The driving forces of these molecular networks are protein interactions rather than the individual functions of single protein components (Pawson, 2004). Biological processes, such as cellular organization, communication, immune responses, and the regulation of transcription and translation, function appropriately only when various proteins interact and work together properly.

Laboratory identification and validation of protein interactions often relies on expensive, time-consuming biochemical and biophysical assays, including ELISA, western blot, immunoprecipitation, Förster resonance energy transfer (FRET) or cross-linking approaches, fluorescence anisotropy, microscale thermophoresis (MST), surface plasmon resonance spectroscopy (SRP), high-throughput screening methods (e.g., phage display), or combinations thereof.

A reliable *in silico* method for predicting protein-protein interactions (PPI) would therefore shed more light on the details of biological pathways and pharmacological responses. Computations may complement and guide biochemical assays. However, explicit molecular dynamics or docking approaches require structural detail, which is often unavailable. Even if the structure of a protein is known, these methods are computationally expensive and therefore impractical for scanning huge libraries of candidates. Therefore, when structural insight is lacking or when speed is crucial, other methods are needed. Non-structure-based computational approaches for identifying potential PPIs generally use an extensive dataset of known protein-protein interactions, combined with information about cellular localization, amino acid sequences, or secondary structures.

These methods may include phylogenetic trees (Pazos and Valencia, 2001), phylogenetic profiles (Barker and Pagel, 2005; Hamp and Rost, 2015a), graph based approaches (Yang *et al.*, 2020), support vector machines (Li *et al.*, 2019), or network-based approaches (Yook *et al.*, 2004; Clauset *et al.*, 2008), stacked autoencoders (Sun *et al.*, 2017), as well as (recurrent) convolutional neural networks (Hashemifar *et al.*, 2018; Chen *et al.*, 2019).

In recent years, distinct prediction methods have been combined, e.g. convolutional neural networks and feature-selecting rotation forests (Wang *et al.*, 2019). Algorithms from language encoding (Yao *et al.*, 2019) and principle component analysis (Kong *et al.*, 2020) were used to derive feature vectors. Structural features were exploited: (Singh *et al.*, 2010; Das and Chakrabarti, 2021). Overall, the field of biology has recently seen a massive increase in applications using deep learning (Ching *et al.*, 2018). Nevertheless, different proteome-wide prediction methods have demonstrated that knowledge of the amino acid sequence alone may be sufficient to identify novel, functional protein-protein interactions (Martin *et al.*, 2005; Shen *et al.*, 2007). These methods usually rely on statistical learning algorithms. Due to its significant advantages, which include simplicity, rapidity, and generality, this method of prediction has become more and more common in recent years (Ofra and Rost, 2003; Betel *et al.*, 2007; Liu *et al.*, 2012; Perovic *et al.*, 2017; Pan *et al.*, 2010; Chen *et al.*, 2020; Das and Chakrabarti, 2021). Precalculated databases are available online: PrePPI (Zhang *et al.*, 2013), ProfPPI (Tran *et al.*, 2018), STRING (Szklarczyk *et al.*, 2011), or PIPS (McDowall *et al.*, 2009). Several webservices offer a limited number of pairwise predictions: PSOPIA (Murakami and Mizuguchi, 2014), iLoops (Planas-Iglesias *et al.*, 2013), iFrag (Garcia-Garcia *et al.*, 2017). To our knowledge, no webserver currently offers scanning entire proteomes.

In this paper, we present an online sequence-based approach to predicting PPIs. The tool’s predictive power was boosted through rigorous fine-tuning of the key elements of machine learning, including dataset generation and feature vector design. Auto-correlation of hydrophobicities combined with a random forest (RF) machine-learning algorithm led to maximum accuracy. The quality and speed of this system make it a suitable high-throughput method for scanning sequence libraries. Furthermore, we achieved a comparable accuracy rate using a graph neural network (GNN). Since this approach has a completely different mathematical structure, we provide it as an option for the user on the server. Additionally, a consensus method is available.

Therefore, **ProteinPrompt** (**protein** prediction of **matching partners**) may serve as a reliable tool for identifying potential interaction partners from an entire proteome. It can thus be used to help identify the yet-unknown biological roles of many proteins and may contribute to identifying new therapeutic targets.

2 Materials and Methods

To maximize the system’s predictive power, it is pivotal to optimize all key elements, including the collection of training and testing data, the calculation of feature vectors, and the selection and fine-tuning of the machine-learning algorithm. Many varied approaches to all these steps were tested. Here, we focus on the approaches that resulted in the highest accuracy rates.

2.1 Collecting data points

In order to create comprehensive training and testing data, we tried to collect as many trustworthy PPI annotations as possible. We included data from various sources, such as the **Database of Human Interacting Proteins**¹ (DIP) (Salwinski *et al.*, 2004), the **Human Protein Reference Database**² (HPRD) (Keshava Prasad *et al.*, 2009), the **Protein Database**³ (PDB) (Berman *et al.*, 2000), and the **Negatome Database**⁴ (Blohm *et al.*, 2014). We also included annotations retrieved from the KUPS⁵ server (Chen *et al.*, 2011), which mainly incorporates PPIs from MINT⁶ (Licata *et al.*, 2012) and IntAct⁷ (Orchard *et al.*, 2014). In addition to the negative annotations collected from the **Negatome Database**, the KUPS server generates negative data points based on the the following criteria: (1) the proteins are functionally dissimilar, (2) the proteins are located in different cellular compartments, and (3) the proteins are part of non-interacting domains.

After intense manual curation and after mapping the different names used to describe the same

proteins, we derived a total list of 31,867 distinct human proteins. For this set, 73,681 positive PPIs were collected from the databases mentioned above. We then used **CD-hit** (Li and Godzik, 2006; Fu *et al.*, 2012) to reduce this set to 41,482 positive protein-protein pairs with at most 50% sequence identity. For the negative pairs, we collected over 1.5 million unique protein-protein pairs, from which we randomly selected comparable number of PPIs to the size of the positive dataset, while still maintaining at most 50% sequence identity.

We separated the data into test and training data to estimate the quality of the optimized prediction model with an independent dataset, which was not involved in the training or similar to the data used for training. Our final training dataset contained 36,423 positive and 34,640 negative PPIs, while our testing dataset contained 5059 positive and 4817 negative data points.

A detailed description of the number of individual proteins in each data set can be found in the Supplementary Section 2.1.

To assess the prediction quality of **ProteinPrompt**, we compared our system to several other publicly available programs. As some of them have limited speed or upload capacity, it was not feasible to perform this test using our entire test dataset, which includes more than 10k data points. Instead, we randomly selected 470 positive and 470 negative pairs from our test set.

2.2 Analysis of data sets

In machine learning, data are often considered equal. Under this assumption, data can be divided into training and test data using random selection. This is clearly not the case for protein sequences, as the level of similarity among sequences may vary dramatically. (A comparison may be drawn from face recognition: Protein sequences can differ from one another as much as human faces differ from those of cats or whales.) Park *et al.* (Park and Marcotte, 2012) point out several combinatorial issues, which were further explored by Hamp *et al.* (Hamp and Rost, 2015b). These issues are due to the pairwise nature of the input and should be considered when building the test dataset. It should be noted that our input data are sequence

¹<http://dip.doe-mbi.ucla.edu/>

²<http://www.hprd.org/>

³www.rcsb.org

⁴<http://mips.helmholtz-muenchen.de/proj/ppi/negatome/>

⁵<http://www.itc.ku.edu/chenlab/kups/>

⁶<https://mint.bio.uniroma2.it/>

⁷<https://www.ebi.ac.uk/intact/>

pairs that are identical in nature and therefore symmetric. This section addresses the relation between data points and their separation into test and training sets.

Obviously, it is significantly easier to predict binding partners for a query sequence that is very similar to one of the sequences in the training dataset than to predict binding partners for a sequence that has no similarity to any sequence in the training set. The latter case requires the algorithm to have "understood" some of principles that control the binding, while the former case only requires it to interpolate from known cases. The more predictions are based on actual understanding, the more general the results will be. However, it is very difficult to understand the complex, 3D interactions of macromolecules based on patterns in their sequences. A full understanding of these interactions requires some level of understanding of the folding of the individual proteins as well as their preferences regarding relative orientation.

Our goal was to train a method that was as general as possible, in the sense that it should not specialize in a certain class of human protein sequences. This meant that we needed to minimize redundancies in the datasets. Furthermore, we intended to perform a very rigorous test by reducing similarities between the training and testing data. As pointed out by Park et al., this approach may not lead to the best overall performance. However, it is the most rigorous way to test such a system.

Therefore, we analyzed the redundancies in our datasets by comparing similarities among data points. It should be noted that we were not looking for pairwise sequence similarity but the similarity of any given pair of sequences with any other pair of sequences. All datasets contain sequence pairs that are known to be binders or non-binders. First, we collected BLAST alignments for all the sequences in all datasets. We calculated similarity by dividing the number of identical positions by the length of the sequence. To compare the similarity of a pair A: (A_1, A_2) with pair B: (B_1, B_2) similarities were calculated in two possible combinations: $\text{Sim}(A_1, B_1)$ with $\text{Sim}(A_2, B_2)$ and $\text{Sim}(A_1, B_2)$ with $\text{Sim}(A_2, B_1)$. The combination resulting in a higher total similarity was selected. The 2D histogram in Figure 1 illustrates the low similarity between the binders in the training and test

Figure 1: Redundancy analysis of binders in the training and test datasets. Depicted is a 2D histogram of similarity between sequence pairs known to be binders in the training and test datasets. The yellow area represents sequence pairs with no or very few similarities (below $\sim 30\%$). It has a count of $1.8 \cdot 10^8$. The violet area represents significantly smaller number of occurrences. A detailed description of this analysis is provided in section 2.2. The values of the histogram can be found in Table S1. Note that this figure is identical to the left image in Figure S1.

datasets. Thus, the selected test set represents a difficult challenge for the algorithm, and the accuracy rates reported here can be considered a worst-case estimate.

Equivalent analyses with very similar results were performed on both datasets separately; the results are provided in the supplementary material (Figures S1, S2, S3). The values that are illustrated in the Figures are provided in Tables S1, S2, S3.

2.3 Selection of learning method

Many machine-learning algorithms are available. We used the `caret` R package (Kuhn, 2008) for fine-tuning and comparison. We tested several artificial neural network implementations, support vector machines, and tree-based methods. After determining that the RF approach performed best on our extensive testing and training datasets, we moved to a Python-based implementation to improve performance. A significant advantage of RFs is that overfitting is not a major problem.

We also extensively tested several more recent machine-learning methods, such as convolu-

Figure 2: Overview of the training procedure.

tional neural networks (CNN) and graph neural networks (GNN), which were implemented using the Tensorflow and PyTorch frameworks, respectively. CNNs were dismissed because of significant problems with overfitting. However, the GNN method resulted in comparable prediction quality to that achieved by the RF. Furthermore, the GNN method does not require any external preprocessing and can handle input data of varied sizes and structures. In our implementation, the GNN translates the data into fixed sized feature vectors, followed by a multilayer perceptron (MLP) that is predicting the binding. Feature vector calculation and predictions are trained in a single optimization scheme. Therefore, the GNN based implementation was trained on the raw data profiles.

2.4 Random forest classification

RF is a supervised learning algorithm that uses an ensemble of classification trees (Breiman, 2001). Each classification tree is built through bootstrap aggregation (bagging), a method of random sampling with replacement. From the original dataset, which contains n feature vectors, a random feature vector is selected n times and then copied to the dataset used to construct the tree. The copied feature vector is not removed from the original dataset. Thus, there will be multiple copies of some feature vectors. This means that different feature vectors are assigned varying degrees of importance. We defined a 420-dimension feature vector $F = (x_1, x_2, \dots, x_{420})$ as the input for the RF model.

To construct an individual tree, a small random subset of a fixed size $m \ll 420$ is extracted from the feature vectors at each node. The best split between these m values, which is the split that leads to the highest predictive power, is selected as the condition at each node. Each tree is grown as large as possible without pruning, resulting in low-bias trees. While this results in overfitting for a single decision tree, RF uses a high number of decision trees, which means that the algorithm has low variance. The final number of trees in the forest was set to 750 to optimize run time.

To further enhance the prediction quality of the RF, we also incorporated inverse and reverse PPI representations in our training. Each protein-protein pair (A, B) is therefore also described in the reversed direction (B, A) as well as two combinations with one inversed protein sequence: (A, B_{inv}) and (B, A_{inv}) . For example A : 'ANLMK', and A_{inv} : 'KMLNA'. Among the eight possible protein pair representations, the combination of the four listed representations yielded the best prediction results and was therefore used in our RF model. This quadrupling of the data was not necessary for the GNN implementation.

2.4.1 Feature vector calculation

Feature vector calculation includes extracting and transforming sequence-based information into a numerical vector of a constant size. Therefore, it is essential to extract the properties that direct the protein-protein interactions.

Each amino acid sequence of a protein-protein complex was transformed into a sequence of numerical values representing seven sequence-derived physicochemical properties. These properties are hydrophobicity (Eisenberg *et al.*, 1984; Koehler *et al.*, 2009), hydrophilicity (Hopp and Woods, 1981), the volume of the side chains of amino acids (Krigbaum and Komoriya, 1979), polarity (Grantham, 1974), polarizability (Charton and Charton, 1982), solvent-accessible surface area (Rose *et al.*, 1985), and the net charge index of the side chains of amino acids (Zhou *et al.*, 2006). The properties were calculated for each residue in the sequence. These scales are commonly used for protein recognition (Ding and Dubchak, 2001) and to predict protein interactions (Bock and Gough,

2001, 2003), protein alignment (Stamm *et al.*, 2013), protein structure (Durham *et al.*, 2009), or protein functional families (Cai *et al.*, 2003). These applications suggest that these properties significantly contribute to the stability of protein-protein complexes. Each amino acid scale was normalized as follows:

$$P'_i = (P_i - \bar{P})/\sigma_P \quad (1)$$

where $\bar{P}$ is the mean and σ is the standard deviation of the scale-based descriptor covering 20 amino acids, respectively:

$$\bar{P} = \sum_{i=1}^{20} P_i/20 \quad (2)$$

and

$$\sigma_P = \sqrt{\frac{1}{20} \sum_{i=1}^{20} (P_i - \bar{P})^2} \quad (3)$$

Auto-correlation (AC) was then used to transform the data into appropriate feature vectors as follows:

$$AC_{lag,j} = \frac{\frac{1}{n-lag} \sum_{i=1}^{n-lag} (S_{i,j} - \bar{S}_j)(S_{i+lag,j} - \bar{S}_j)}{\sigma_{S_j} * \sigma_{S_j}} \quad (4)$$

where S_j is the translated amino acid sequence using the normalized scale-based descriptor P'_j with $j = 1, 2, \dots, 7$, n is the length of sequence S , $lag = 1, 2, \dots, 30$ is the shift for which the auto-correlation is calculated, $\bar{S}_j$ and σ_{S_j} are the mean and standard deviation of the translated sequence, respectively. Ding *et al.* (2016) showed that a maximum lag of less than 30 tends to lose useful information, while larger values may induce noise. Accordingly, the number of AC values for each of the seven scales is 30. The feature vector describing any individual amino acid sequence has $7 \cdot 30 = 210$ elements or dimensions. Thus, for a pair of sequences, the feature space has 420 dimensions.

2.5 Graph neural networks

Graph neural networks (GNN) are a relatively new type of neural networks that operate on graphs (Scarselli *et al.*, 2009), (Battaglia *et al.*, 2018). Features can be assigned and predicted on a node, edge, or graph level. The algorithm takes a graph

as input, performs computations on the graph itself through a process called message passing, and returns a graph of identical structure but updated features.

In a message passing round, first, the edge features are updated by a learnable update function that takes connected nodes and current edge features into account. Afterwards, another learnable function calculates new node features based on the aggregation of all connected edges and current node features. In the final step, a graph-level target is constructed by a third learnable function that takes an aggregation of all the nodes and edges as input. Through this process, every node and edge collects information about its local region in the graph; these data are used to infer global features.

We used GNNs to condense the information in protein sequences of varying lengths to a fixed-sized vector representation of each sequence. One graph is constructed for each sequence; nodes represent amino acids, and edges are constructed between the nodes of adjacent amino acids. The node features were encoded using a MLP that transformed the values of the amino acid profiles into a 32-dimension vector. The edge features were encoded from a vector of same size that was initialized with ones. The update functions for the nodes and edges and for the graph were implemented by MLPs.

After five rounds of message passing, a final aggregation function is performed on the graph to calculate a 128-dimension graph-level feature from all nodes and edges. This output can be understood as an abstract representation of each protein. This is done for both sequences. Then, their feature vectors are concatenated to one 256-dimension vector, which is then used as the input for the MLP. The last network returns the prediction of the binding propensity of the two proteins.

The model was built using PyTorch and the PyTorch geometric library. The message passing graph net block was realized using its MetaLayer class with an additional edge update on the graph. Parameters were initialized as described by (Glorot and Bengio, 2010), and a leaky rectified linear unit (ReLU) (Maas *et al.*, 2013) was used for activations, except in the final layer of the last MLP, where a sigmoid activation was used. Training

was done using binary cross-entropy loss and the AdamW (Loshchilov and Hutter, 2017) optimizer with a learning rate of $1e-5$ for 350 epochs. The iteration with the best performance on the test set was refined over an additional 50 epochs with a learning rate of $1e-7$. See Figure S4 in the supplementary materials for a graphical illustration of our implementation.

2.6 Consensus of RF and GNN

We trained a simple neural network with the outputs of RF and GNN for performing a consensus prediction. The scores returned by RF and GNN show rather different distributions. Thus, a linear combination is unlikely to result in the best possible consensus. We selected a NN with two inputs, two hidden layers with 8 neurons and a single output neuron (2-8-8-1).

3 Results

3.1 Performance RF

Based on our test dataset of 5094 positive and 5094 negative cases (binders vs. non-binders), we plotted the receiver operation characteristics (ROC) curve, shown in Figure 3. This allowed us to estimate the overall quality of the predictions and also provides a visual overview of the relationship between true and false positives. Each point on the curve shows how many falsely predicted binders should be expected for a given number of correctly predicted ones. As the datapoints are sorted by their predicted binding propensity, a reasonable threshold can be selected. This allows to fine-tune the balance between sensitivity and specificity.

Sensitivity is the ratio of correctly predicted binders (true positives) to actual binders. Specificity is the ratio of correctly predicted non-binders (true negatives) to all non-binders. For example, in Figure 3, a false positive rate of 0.03 will result in a true positive rate of ~ 0.6 . Thus, selecting this point (its associated score) as the threshold will result in a hit rate of 60% and only 3% false predictions should be expected.

An ideal signal would lead to a rectangular plot with an area under the curve (AUC) of 1. The other extreme, pure noise, would result in a diagonal line with an AUC of 0.5. Our method results

Figure 3: Performance of ProteinPrompt on the entire test dataset of over 10,000 PPIs.

Tool	AUC	Spec.	Sens.	Acc.
ProteinPrompt	0.94	0.88	0.84	0.86
SPPS	0.77	0.34	0.96	0.66
TRI.tool	0.73	0.95	0.32	0.63
LR.PPI	0.60	0.16	0.91	0.53

Table 1: Quality measures of ProteinPrompt compared to other publicly available tools, such as SPPS (Liu *et al.*, 2012), TRI.tool (Perovic *et al.*, 2017), and LR.PPI (Pan *et al.*, 2010). AUC, specificity, sensitivity, and accuracy are listed.

in an AUC of 0.95, specificity of 0.88, sensitivity of 0.87, and an accuracy rate of 0.88. Currently, our method balances specificity and sensitivity to avoid unwanted bias in different applications.

3.1.1 RF compared to other tools

We compared our optimized RF system to publicly available tools such as SPPS (Liu *et al.*, 2012), TRI.tool (Perovic *et al.*, 2017), and LR.PPI (Pan *et al.*, 2010). Due to limited access, we used the reduced test dataset, as described in the methods section. We also re-evaluated ProteinPrompt on the reduced test dataset to ensure comparable results. The results of our evaluation using the reduced test dataset are somewhat different from those we obtained using the entire test dataset.

The results shown in Table 1 and the ROC

Figure 4: ROC curves for ProteinPrompt, SPPS, TRI_tool, LR_PPI using the reduced test dataset of 968 protein-protein pairs.

plot in Figure 4 suggest that ProteinPrompt outperforms the other three methods. Furthermore, ProteinPrompt balances sensitivity and specificity. SPPS and LR_PPI have excellent sensitivity of 0.96 and 0.91; however, their specificity is rather poor: 0.34 and 0.16, respectively. TRI_tool on the other hand, shows a massive bias towards specificity, which is 0.95, compared to a sensitivity of 0.32. Furthermore, the overall prediction accuracy and the AUC of our tool are significantly higher. Accuracy: 0.86 versus 0.66 (SPPS), 0.63 (TRI_tool) or 0.53 (LR_PPI). AUC: 0.94 versus 0.77 (SPPS), 0.73 (TRI_tool), and 0.60 (LR_PPI), respectively.

3.1.2 Individual test cases

We further tested the detection rate of ProteinPrompt using experimentally verified protein interaction partners (EV PPIs). We compared the output of ProteinPrompt to that of the STRING database⁸, which has been shown to include a very high number of experimentally proven PPIs (Bajpai *et al.*, 2020). Here, the EV PPIs of five different prominent proteins with various cellular functions (Table S4 in the supplementary material) with high confidence (score ≥ 0.7) were investigated. The PPI predictions of ProteinPrompt were compared to those of the

⁸<https://string-db.org>

Figure 5: EV PPI scores from ProteinPrompt (green) and from STRING database (grey) for a test set of five different proteins, indicated by their gene names in the boxplot.

STRING database. ProteinPrompt found all but one of the EV PPIs output by the STRING Database; for the SRC gene one of nine binding partners was not identified (Table S4). The scores given by the STRING database and the scores of ProteinPrompt were then statistically analyzed and plotted as boxplots to enable direct comparison (Figure 5). On average, ProteinPrompt predicted all identified EV PPIs with comparable or better accuracy than the STRING database (Figure 5, Table S4). However, due to the complexity of PPIs in nature, STRING scores do not necessarily reflect real protein-protein affinities. On our test dataset, ProteinPrompt was able to rapidly predict EV PPIs with an average accuracy rate of 0.89 and a SD of 0.09. From a user's perspective, ProteinPrompt identified 98% of all EV PPIs with a predicted score above 0.7, 84% of all EV PPIs with a predicted score above 0.8, and 46% of all EV PPIs with a predicted score above 0.9.

3.1.3 Dataset of Park and Marcotte

As an additional test, we used the Park and Marcotte (2012) dataset on which a number of methods were tested, allowing a direct comparison. Park and Marcotte (2012) tested 7 methods, Ding *et al.* (2016) extended this with an additional 3 methods. ProteinPrompt in its current form performs similarly on this dataset as it does on our own test data. For all human subsets, ProteinPrompt achieved significantly higher accuracy than all other meth-

ods. When ProteinPrompt was trained with this data, the performance is similar to the other top ranking tools. For most human subsets, the re-trained version of ProteinPrompt ranked second. First, this confirms that combining RF with autocorrelation has the potential to achieve high accuracy. Second, this test highlights the critical impact of the dataset on the final performance. For the yeast data, ProteinPrompt does not perform as well, indicating limited transferability to other organisms, a fact also observed with other learning methods. Most learning methods are only valid within the domain of their training data. A detailed description can be found in the Supplementary Materials.

3.2 Performance GNN

In the first test, we used the same seven amino acid scales that were used for the RF. We compared this to a graph topology in which 13 and 15 scales were assigned to each amino acid. The resulting accuracy rates are listed in Table 2.

# profiles	Accuracy
7	83.5%
13	86.0%
15	83.8%

Table 2: Accuracies of GNNs trained on different numbers of amino acid profiles.

Adding as many residue scales as possible to the graph does not improve the accuracy. Currently, the maximum accuracy is 86% for 13 residue scales. Extensive future research would be needed to further maximize the performance of the GNN.

In cross validations, in which 20% of the training data were used as validation set, the accuracies for the validation data were nearly identical to the one of the test data. This reflects the low level of redundancy within the training dataset, which is shown in Figure S3. We therefore considered it acceptable to perform a final optimization with 95% of the test dataset added to the training data. The resulting model is the one uploaded to the server. This strategy should lead to the most generally applicable model.

3.3 Performance consensus RF and GNN

The consensus using the RF and GNN as input was leading to a slight improvement with a final accuracy of 89.4%.

3.4 Webserver implementation

ProteinPrompt is available online as a webserver at:

<http://proteinformatics.org/ProteinPrompt>
Basic usage is as simple as providing an input sequence. By default, ProteinPrompt will search our manually curated database of human proteins, which contains 27,223 sequences. Scanning the entire human database takes approximately one minute for the default method RF. For GNN the search is even faster. The consensus has to execute both methods and therefore requires correspondingly more time. Other, more extensive databases are also provided, including mammalian, vertebrate, and metazoan protein sets. Searching these databases takes considerably longer. For example, the vertebrate database has 91,592 entries; therefore, searching it takes approximately three times longer than searching the human protein database.

The server is free for academic users. Providing an email address is optional. ProteinPrompt was originally optimized for sequences with a minimum length of 16 AA. The sequence length of the uploaded proteins is detected automatically. When the user provides a sequence shorter than 16 AA, a warning appears, but the calculation is still performed. As output, a ranked and scored list of proteins from the database is returned; this list can be downloaded.

4 Discussion

ProteinPrompt offers a reliable, fast way to predict protein interaction partners based on protein sequences. It is available as an easy-to-use online tool and is thus accessible to non-expert users. In order to develop this fast, reliable service, we optimized the learning algorithm, the binding database, and the representations of the sequences. We determined that the RF algorithm combined with autocorrelation on seven amino acid scales resulted in the highest accuracy.

We also determined that the GNN method performed nearly as well as the RF algorithm. To support consensus building, we offer both implementations on our server. It is remarkable that the RF algorithm, which is conceptually comparably simple, performs so well on such a complex task. This is even more remarkable considering that the RF approach, unlike the GNN, does not include simultaneous optimization of feature vector creation and model building.

An extensive database with limited redundancy was essential to reliably test the tool. This database was obtained through several iterations of manual curation of the test and training datasets. Despite strict separation of the training and test datasets, which posed a significant learning challenge, **ProteinPrompt** turned out to perform very well compared to other available servers and methods. **ProteinPrompt** is reasonably accurate and can scan the entire human proteome within approximately one minute. To our knowledge, **ProteinPrompt** is currently the fastest online service available for scanning different proteomes to identify potential binding partners based on a sequence level. It is reasonable to assume that expanding the training data would lead to higher accuracy rates. Based on our extensive tests, we expect **ProteinPrompt** to support a better understanding of the complex networks of protein-protein interactions, which are the basis for a broad range of biological mechanisms.

Acknowledgements

We extend our thanks to Johanna Tiemann, Alexander Vogel, Thorsten Kaiser and Daniel Wiegrefe for proofreading.

Funding

This work was funded by the Sächsische Aufbaubank (SAB) [100256197, 100267985], and Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC 1423 [421152132; sub- projects Z04]

References

- Bajpai, A. K., Davuluri, S., Tiwary, K., Narayanan, S., Oguru, S., Basavaraju, K., Dayalan, D., Thirumurugan, K., and Acharya, K. K. (2020). Systematic comparison of the protein-protein interaction databases from a user’s perspective. *Journal of Biomedical Informatics*, **103**, 103380.
- Barker, D. and Pagel, M. (2005). Predicting functional gene links from phylogenetic-statistical analyses of whole genomes. *PLoS Comput Biol*, **1**(1), e3.
- Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., Tacchetti, A., Raposo, D., Santoro, A., Faulkner, R., Gulcehre, C., Song, F., Ballard, A., Gilmer, J., Dahl, G., Vaswani, A., Allen, K., Nash, C., Langston, V., Dyer, C., Heess, N., Wierstra, D., Kohli, P., Botvinick, M., Vinyals, O., Li, Y., and Pascanu, R. (2018). Relational inductive biases, deep learning, and graph networks.
- Berman, H. M., Westbrook, J., Feng, Z., Gilliland, G., Bhat, T. N., Weissig, H., Shindyalov, I. N., and Bourne, P. E. (2000). The protein data bank. *Nucleic Acids Res*, **28**(1), 235–42.
- Betel, D., Breitkreuz, K. E., Isserlin, R., Dewar-Darch, D., Tyers, M., and Hogue, C. W. (2007). Structure-templated predictions of novel protein interactions from sequence information. *PLoS Comput Biol*, **3**(9), 1783–9.
- Blohm, P., Frishman, G., Smialowski, P., Goebels, F., Wachinger, B., Ruepp, A., and Frishman, D. (2014). Negatome 2.0: a database of non-interacting proteins derived by literature mining, manual annotation and protein structure analysis. *Nucleic Acids Res*, **42**(Database issue), D396–400.
- Bock, J. R. and Gough, D. A. (2001). Predicting protein-protein interactions from primary structure. *Bioinformatics*, **17**(5), 455–60.
- Bock, J. R. and Gough, D. A. (2003). Whole-proteome interaction mining. *Bioinformatics*, **19**(1), 125–34.
- Breiman, L. (2001). Random forests. *Machine Learning*, **45**(1), 5–32.
- Cai, C. Z., Han, L. Y., Ji, Z. L., Chen, X., and Chen, Y. Z. (2003). Svm-prot: Web-based support vector machine software for functional classification of a protein from its primary sequence. *Nucleic Acids Res*, **31**(13), 3692–7.
- Charton, M. and Charton, B. I. (1982). The structural dependence of amino acid hydrophobicity parameters. *J Theor Biol*, **99**(4), 629–44.
- Chen, M., Ju, C. J., Zhou, G., Chen, X., Zhang, T., Chang, K. W., Zaniolo, C., and Wang, W. (2019). Multifaceted protein-protein interaction prediction based on Siamese residual RCNN. *Bioinformatics*, **35**(14), i305–i314.

- Chen, X. W., Jeong, J. C., and Derymyer, P. (2011). Kups: constructing datasets of interacting and non-interacting protein pairs with associated attributions. *Nucleic Acids Res*, **39**(Database issue), D750–4.
- Chen, Y., Wang, W., Liu, J., Feng, J., and Gong, X. (2020). Protein Interface Complementarity and Gene Duplication Improve Link Prediction of Protein-Protein Interaction Network. *Front Genet*, **11**, 291.
- Ching, T., Himmelstein, D. S., Beaulieu-Jones, B. K., Kalinin, A. A., Do, B. T., Way, G. P., Ferrero, E., Agapow, P. M., Zietz, M., Hoffman, M. M., Xie, W., Rosen, G. L., Lengerich, B. J., Israeli, J., Lanchantin, J., Woloszynek, S., Carpenter, A. E., Shrikumar, A., Xu, J., Cofer, E. M., Lavender, C. A., Turaga, S. C., Alexandari, A. M., Lu, Z., Harris, D. J., DeCaprio, D., Qi, Y., Kundaje, A., Peng, Y., Wiley, L. K., Segler, M., Boca, S. M., Swamidass, S. J., Huang, A., Gitter, A., and Greene, C. S. (2018). Opportunities and obstacles for deep learning in biology and medicine. *J R Soc Interface*, **15**(141).
- Clauset, A., Moore, C., and Newman, M. E. (2008). Hierarchical structure and the prediction of missing links in networks. *Nature*, **453**(7191), 98–101.
- Das, S. and Chakrabarti, S. (2021). Classification and prediction of protein-protein interaction interface using machine learning algorithm. *Sci Rep*, **11**(1), 1761.
- Ding, C. H. and Dubchak, I. (2001). Multi-class protein fold recognition using support vector machines and neural networks. *Bioinformatics*, **17**(4), 349–58.
- Ding, Y., Tang, J., and Guo, F. (2016). Predicting protein-protein interactions via multivariate mutual information of protein sequences. *BMC Bioinformatics*, **17**(1), 398.
- Durham, E., Dorr, B., Woetzel, N., Staritzbichler, R., and Meiler, J. (2009). Solvent accessible surface area approximations for rapid and accurate protein structure prediction. *J Mol Model*, **15**(9), 1093–108.
- Eisenberg, D., Schwarz, E., Komaromy, M., and Wall, R. (1984). Analysis of membrane and surface protein sequences with the hydrophobic moment plot. *J Mol Biol*, **179**(1), 125–42.
- Fu, L., Niu, B., Zhu, Z., Wu, S., and Li, W. (2012). Cd-hit: accelerated for clustering the next-generation sequencing data. *Bioinformatics*, **28**(23), 3150–2.
- Garcia-Garcia, J., Valls-Comamala, V., Guney, E., Andreu, D., Muoz, F. J., Fernandez-Fuentes, N., and Oliva, B. (2017). iFrag: A Protein-Protein Interface Prediction Server Based on Sequence Fragments. *J Mol Biol*, **429**(3), 382–389.
- Glorot, X. and Bengio, Y. (2010). Understanding the difficulty of training deep feedforward neural networks. In Y. W. Teh and M. Titterton, editors. *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy. PMLR.
- Grantham, R. (1974). Amino acid difference formula to help explain protein evolution. *Science*, **185**(4154), 862–4.
- Hamp, T. and Rost, B. (2015a). Evolutionary profiles improve protein-protein interaction prediction from sequence. *Bioinformatics*, **31**(12), 1945–1950.
- Hamp, T. and Rost, B. (2015b). More challenges for machine-learning protein interactions. *Bioinformatics*, **31**(10), 1521–1525.
- Hashemifar, S., Neyshabur, B., Khan, A. A., and Xu, J. (2018). Predicting protein-protein interactions through sequence-based deep learning. *Bioinformatics*, **34**(17), i802–i810.
- Hopp, T. P. and Woods, K. R. (1981). Prediction of protein antigenic determinants from amino acid sequences. *Proc Natl Acad Sci U S A*, **78**(6), 3824–8.
- Keshava Prasad, T. S., Goel, R., Kandasamy, K., Keerthikumar, S., Kumar, S., Mathivanan, S., Telikicherla, D., Raju, R., Shafreen, B., Venugopal, A., Balakrishnan, L., Marimuthu, A., Banerjee, S., Somanathan, D. S., Sebastian, A., Rani, S., Ray, S., Harrys Kishore, C. J., Kanth, S., Ahmed, M., Kashyap, M. K., Mohmood, R., Ramachandra, Y. L., Krishna, V., Rahiman, B. A., Mohan, S., Ranganathan, P., Ramabadran, S., Chaerkady, R., and Pandey, A. (2009). Human protein reference database–2009 update. *Nucleic Acids Res*, **37**(Database issue), D767–72.
- Koehler, J., Woetzel, N., Staritzbichler, R., Sanders, C. R., and Meiler, J. (2009). A unified hydrophobicity scale for multispan membrane proteins. *Proteins*, **76**(1), 13–29.
- Kong, M., Zhang, Y., Xu, D., Chen, W., and Dehmer, M. (2020). Weighted Sparse Representation Based Classification. *Front Genet*, **11**, 18.
- Krigbaum, W. R. and Komoriya, A. (1979). Local interactions as a structure determinant for protein molecules: II. *Biochim Biophys Acta*, **576**(1), 204–48.
- Kuhn, M. (2008). Building predictive models in r using the caret package. *Journal of Statistical Software, Articles*, **28**(5), 1–26.
- Li, W. and Godzik, A. (2006). Cd-hit: a fast program for clustering and comparing large sets of protein or nucleotide sequences. *Bioinformatics*, **22**(13), 1658–9.
- Li, X., Yang, L., Zhang, X., and Jiao, X. (2019). Prediction of Protein-Protein Interactions Based on Domain. *Comput Math Methods Med*. **2019**, 5238406.

- Licata, L., Briganti, L., Peluso, D., Perfetto, L., Iannuccelli, M., Galeota, E., Sacco, F., Palma, A., Nardoza, A. P., Santonico, E., Castagnoli, L., and Cesareni, G. (2012). Mint, the molecular interaction database: 2012 update. *Nucleic Acids Res*, **40**(Database issue), D857–61.
- Liu, X., Liu, B., Huang, Z., Shi, T., Chen, Y., and Zhang, J. (2012). Spps: a sequence-based method for predicting probability of protein-protein interaction partners. *PLoS One*, **7**(1), e30938.
- Loshchilov, I. and Hutter, F. (2017). Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- Maas, A. L., Hannun, A. Y., Ng, A. Y., et al. (2013). Rectifier nonlinearities improve neural network acoustic models. In *Proc. icml*, volume 30, page 3. Citeseer.
- Martin, S., Roe, D., and Faulon, J. L. (2005). Predicting protein-protein interactions using signature products. *Bioinformatics*, **21**(2), 218–26.
- McDowall, M. D., Scott, M. S., and Barton, G. J. (2009). Pips: human protein-protein interaction prediction database. *Nucleic Acids Res*, **37**(Database issue), D651–6.
- Murakami, Y. and Mizuguchi, K. (2014). Homology-based prediction of interactions between proteins using Averaged One-Dependence Estimators. *BMC Bioinformatics*, **15**, 213.
- Ofran, Y. and Rost, B. (2003). Predicted protein-protein interaction sites from local sequence information. *FEBS Lett*, **544**(1-3), 236–9.
- Orchard, S., Ammari, M., Aranda, B., Breuza, L., Briganti, L., Broackes-Carter, F., Campbell, N. H., Chavali, G., Chen, C., del Toro, N., Duesbury, M., Dumousseau, M., Galeota, E., Hinz, U., Iannuccelli, M., Jagannathan, S., Jimenez, R., Khadake, J., Lagreid, A., Licata, L., Lovering, R. C., Meldal, B., Melidoni, A. N., Milagros, M., Peluso, D., Perfetto, L., Porras, P., Raghunath, A., Ricard-Blum, S., Roechert, B., Stutz, A., Tognolli, M., van Roey, K., Cesareni, G., and Hermjakob, H. (2014). The mintact project—intact as a common curation platform for 11 molecular interaction databases. *Nucleic Acids Res*, **42**(Database issue), D358–63.
- Pan, X. Y., Zhang, Y. N., and Shen, H. B. (2010). Large-scale prediction of human protein-protein interactions from amino acid sequence based on latent topic features. *J Proteome Res*, **9**(10), 4992–5001.
- Park, Y. and Marcotte, E. M. (2012). Flaws in evaluation schemes for pair-input computational predictions. *Nat. Methods*, **9**(12), 1134–1136.
- Pawson, T. (2004). Specificity in signal transduction: from phosphotyrosine-sh2 domain interactions to complex cellular systems. *Cell*, **116**(2), 191–203.
- Pazos, F. and Valencia, A. (2001). Similarity of phylogenetic trees as indicator of protein-protein interaction. *Protein Eng*, **14**(9), 609–14.
- Perovic, V., Sumonja, N., Gemovic, B., Toska, E., Roberts, S. G., and Veljkovic, N. (2017). Tri-tool: a web-tool for prediction of protein-protein interactions in human transcriptional regulation. *Bioinformatics*, **33**(2), 289–291.
- Planas-Iglesias, J., Marin-Lopez, M. A., Bonet, J., Garcia-Garcia, J., and Oliva, B. (2013). iLoops: a protein-protein interaction prediction server based on structural features. *Bioinformatics*, **29**(18), 2360–2362.
- Rose, G. D., Geselowitz, A. R., Lesser, G. J., Lee, R. H., and Zehfus, M. H. (1985). Hydrophobicity of amino acid residues in globular proteins. *Science*, **229**(4716), 834–8.
- Salwinski, L., Miller, C. S., Smith, A. J., Pettit, F. K., Bowie, J. U., and Eisenberg, D. (2004). The database of interacting proteins: 2004 update. *Nucleic Acids Res*, **32**(Database issue), D449–51.
- Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., and Monfardini, G. (2009). The graph neural network model. *IEEE Trans Neural Netw*, **20**(1), 61–80.
- Shen, J., Zhang, J., Luo, X., Zhu, W., Yu, K., Chen, K., Li, Y., and Jiang, H. (2007). Predicting protein-protein interactions based only on sequences information. *Proc Natl Acad Sci U S A*, **104**(11), 4337–41.
- Singh, R., Park, D., Xu, J., Hosur, R., and Berger, B. (2010). Struct2Net: a web service to predict protein-protein interactions using a structure-based approach. *Nucleic Acids Res*, **38**(Web Server issue), W508–515.
- Stamm, M., Staritzbichler, R., Khafizov, K., and Forrest, L. R. (2013). Alignment of helical membrane protein sequences using alignme. *PLoS One*, **8**(3), e57731.
- Sun, T., Zhou, B., Lai, L., and Pei, J. (2017). Sequence-based prediction of protein protein interaction using a deep-learning algorithm. *BMC Bioinformatics*, **18**(1), 277.
- Szklarczyk, D., Franceschini, A., Kuhn, M., Simonovic, M., Roth, A., Minguez, P., Doerks, T., Stark, M., Muller, J., Bork, P., Jensen, L. J., and von Mering, C. (2011). The string database in 2011: functional interaction networks of proteins, globally integrated and scored. *Nucleic Acids Res*, **39**(Database issue), D561–8.
- Tran, L., Hamp, T., and Rost, B. (2018). ProfPPIdb: Pairs of physical protein-protein interactions predicted for entire proteomes. *PLoS One*, **13**(7), e0199988.
- Wang, L., Wang, H. F., Liu, S. R., Yan, X., and Song, K. J. (2019). Predicting Protein-Protein Interactions from Matrix-Based Protein Sequence Using Convolution Neural Network and Feature-Selective Rotation Forest. *Sci Rep*, **9**(1), 9848.

- Yang, F., Fan, K., Song, D., and Lin, H. (2020). Graph-based prediction of Protein-protein interactions with attributed signed graph embedding. *BMC Bioinformatics*, **21**(1), 323.
- Yao, Y., Du, X., Diao, Y., and Zhu, H. (2019). An integration of deep learning with feature embedding for protein-protein interaction prediction. *PeerJ*, **7**, e7126.
- Yook, S. H., Oltvai, Z. N., and Barabási, A. L. (2004). Functional and topological characterization of protein interaction networks. *Proteomics*, **4**(4), 928–42.
- Zhang, Q. C., Petrey, D., Garzás, J. I., Deng, L., and Honig, B. (2013). PrePPI: a structure-informed database of protein-protein interactions. *Nucleic Acids Res*, **41**(Database issue), D828–833.
- Zhou, P., Tian, F.-F., Li, B., Wu, S.-R., and Li, Z.-L. (2006). Genetic algorithm-based virtual screening of combinative mode for peptide/protein. **64**, 691–697.