Multi-Task Learning to Leverage Partially Annotated Data for PPI Interface Prediction

Henriette Capel

IBIVU – Center for Integrative Bioinformatics, Vrije Universiteit Amsterdam

K. Anton Feenstra

IBIVU – Center for Integrative Bioinformatics, Vrije Universiteit Amsterdam

Sanne Abeln (✉ s.abeln@vu.nl)

IBIVU – Center for Integrative Bioinformatics, Vrije Universiteit Amsterdam

Research Article

Keywords:

Posted Date: February 8th, 2022

DOI: <https://doi.org/10.21203/rs.3.rs-1269779/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Multi-task learning to leverage partially annotated data for PPI interface prediction

Henriette Capel¹, K. Anton Feenstra¹, and Sanne Abeln^{1,*}

¹IBIVU – Center for Integrative Bioinformatics, Vrije Universiteit Amsterdam, Amsterdam 1081HV, The Netherlands

*s.abeln@vu.nl

ABSTRACT

Protein protein interactions (PPI) are crucial for protein functioning, nevertheless predicting residues in PPI interfaces from the protein sequence remains a challenging problem. In addition, structure-based functional annotations, such as the PPI interface annotations, are scarce: only for about one-third of all protein structures residue-based PPI interface annotations are available. If we want to use a deep learning strategy, we have to overcome the problem of limited data availability. Here we use a multi-task learning strategy that can handle missing data. We start with the multi-task model architecture, and adapted it to carefully handle missing data in the cost function. As related learning tasks we include prediction of secondary structure, solvent accessibility and buried residue. Our results show that the multi-task learning strategy significantly outperforms single task approaches. Moreover, only the multi-task strategy is able to effectively learn over a dataset extended with structural feature data, without additional PPI annotations. The multi-task setup becomes even more important, if the fraction of PPI annotations becomes very small: the multi-task learner trained on only one-eighth of the PPI annotations – with data extension – reaches the same performances as the single-task learner on all PPI annotations. Thus, we show that the multi-task learning strategy can be beneficial for a small training dataset where the protein's functional properties of interest are only partially annotated.

Introduction

Protein sequence databases¹ continue to grow rapidly and structural information is becoming more readily available². Nevertheless, precise functional annotation based on the protein structure, such as protein binding sites³, are still scarce, and difficult to predict. Therefore, computational techniques are used to predict several functional structural properties of proteins based on the protein sequence. One of these properties is the physical interaction interface between proteins which are crucial for the functioning of a protein⁴. Interaction between proteins is required in many biological processes, such as DNA replication, RNA transcription, signal transduction, control of cellular processes, protein transport and metabolism⁵⁻⁹. Furthermore, many diseases can be related to the deformation of a protein's interface^{10,11}. Predicting the set of residues in a protein that interact with other proteins is an important, but still challenging task¹². Moreover, structural information on residues that make up the interface is scarce. The size of the PPI annotated database is only a small fraction of the size of structural annotated database. The size of the structural annotated database, in turn, is a small fraction of the size of the protein sequence database (see [Figure 1](#)). To efficiently train deep neural networks for the PPI interface prediction task we have to overcome the problem of the limited size of the training dataset.

Due to the successes of deep learning in fields such as natural language processing, deep learning approaches are increasingly used and have shown great successes for protein structural feature prediction¹³⁻¹⁶. In deep learning, multiple connected layers, along with their parameters, predict the output of the corresponding input features¹⁷. Approaches and models such as convolutional neural

Figure 1. Comparison of the number of entries available in databases on protein sequence, protein structure and specific structure-based functional annotations: protein-protein interaction (PPI) interface. These results are respectively based on the protein entries available in the UniProtKB/TrEMBL database, protein entries available in the Protein Data Bank (PDB), and the protein entries with PPI interface annotations. Note that the y-axis is logarithmic.

networks (CNN), residual neural networks (ResNet), recurrent neural networks (RNN), long short term memory networks (LSTM), transformers and multi-task learners appear in recent structure prediction methods^{14,15,18–21}. Hanson et al.¹⁵ used, among others, ultra-deep ResNets in the SPOT-1D model which were able to capture non-local interactions between residues that are only close in the protein structure and not in the protein sequence¹⁵. Heffernan et al.²⁰ used LSTM bidirectional RNNs and showed that this method is useful to capture long range interactions, especially for residues with large numbers of long-range contacts. We recently compared the usage of different Neural network architectures for the prediction of protein interfaces²². Furthermore, transformers have been successfully used in the language of proteins^{23,24}. In transformers, information learned from general domain data, like protein sequences, is transferred to domain specific data, such as secondary structure prediction. Another strategy in which information is transferred is multi-task learning.

In contrast to single-task learning, in which the aim is to improve the performance of one specific prediction task, for multi-task learning²⁵ the aim is to improve the performance of multiple learning tasks simultaneously. Training the multi-task model on different tasks at the same time allows to the model to learn a shared representation, providing a way to transfer information learned between specific tasks^{13,25}. Multi-task learning is related to inductive transfer learning²⁶. The main difference is that the aim of inductive transfer learning is to only achieve high performances for the main task, whereas the aim of multi-task learning is to learn both the main and related tasks²⁶. The multi-task learning strategy can be implemented in the end-to-end learning architecture of deep learning models. In order to learn the model to make accurate predictions for all tasks, the loss of the various tasks should be represented in the loss function used during training the multi-task model¹³. This strategy has previously been applied to the domain of proteins by, for example, concurrently predicting multiple protein structural properties¹⁴. The benefit of using information stored in a related protein structural prediction task for other protein structural prediction tasks, among which PPI interface prediction, has already been shown for single-task learners which use the related tasks as input features^{7,27,28}. One advantage of the multi-task strategy is not having to generate input features a priori when applying the model to

a new input. Here we are interested both in the inductive transfer's learning ability of the multi-task setup to improve the performance of PPI interface prediction, as well as the actual predictions of the related tasks, such as surface accessibility, as these may provide, for example, insight in the nature of the binding site.

Protein structural property prediction methods commonly use a multi-task learning strategy. Note that the labels for many structural properties, such as secondary structure and solvent accessibility, can be generated only if the three-dimensional structure of a protein is available. Klausen et al.¹⁴ built the sequence-based method NetSurfP-2.0 to predict solvent accessibility, secondary structure, structural disorder and backbone dihedral angles, using a combination of CNNs and LSTMs. The deep learning model SPOT-1D is based on an ensemble of ResNets and CNNs to predict secondary structure, backbone angles, solvent accessibility and contact number¹⁵. Secondary structure prediction by SPOT-1D resulted in higher performances than reached by NetSurfP-2.0¹⁵. In 2020, Xu et al.¹⁶ published their method OPUS-TASS. This multi-task learner, based on a combination of CNNs, transformers and LSTMs, was able to improve the predictions of secondary structure and backbone angles even further¹⁶. Xu et al.¹⁶ trained multiple models, including a different set of the learning tasks: secondary structure in three and eight classes, backbone torsion angles, absolute solvent accessibility, side-chain dihedral angles, and the local backbone structure descriptor CSF3¹⁶. The multi-task learning strategy has, to the best of our knowledge, not been used for PPI interface prediction.

PPI prediction models were recently described in the review paper by Savojardo et al.¹². In this paper the different models were distinguished into methods using the primary protein sequence as input, and methods using the three-dimensional protein structure as input for the prediction model. Moreover, methods can be partner independent or partner dependent¹². In this study, we predict PPI interface residues based on the primary sequence in a partner-unspecific way. The most recent other sequence-based partner-unspecific models are SSWRF²⁹, SeRenDIP^{7,30} and SCRIBER⁹. The SSWRF method uses an ensemble support vector machine and a sample-weighted random forest to predict the PPI interface²⁹. SeRenDIP is a random forest model trained on datasets containing either only homomeric interactions, only heteromeric, or containing both types of interactions⁷. SCRIBER is a model based on multilevel logistic regression and trained on a dataset containing multiple types of protein interactions. In this study, we used the multi-task learning approach for PPI interface prediction using a deep learning model architecture based on the OPUS-TASS model¹⁶. Note that the highest accuracy for protein structural prediction tasks can only be obtained if conservation is used as an input feature³¹. Existing PPI interface prediction methods use conservation^{7-9,29,32-34}, secondary structure^{8,9,27,28}, surface accessibility^{8,9,29,35-37}, backbone flexibility^{38,39} or a combination of these^{7,30} as input features. Previous studies showed that high solvent accessible residues are more likely to be interface residues³⁵⁻³⁷.

In this work, we take sequence-derived properties and sequence-conservation as input features, similar to OPUS-TASS¹⁶, but do not use predicted structural properties such as secondary structure or surface accessibility as input features – as is common in many PPI interface prediction methods. Instead, these structural properties are used as related learning tasks in our multi-task setup, as shown schematically in [Figure 2](#).

Here, we investigate if PPI interface prediction can be improved by defining the task as a multi-task learning problem. We show model performances for different combinations of the related learning tasks: secondary structure in three and eight classes, absolute solvent accessibility, and buried residues. Using these tasks, higher prediction performances are reached compared to the single-task interface prediction model. In addition, the multi-task setup offers the possibility of training on a partially annotated dataset by continuing learning on exclusively the related tasks. The dataset used in this study is only partially annotated with PPI interface labels. The benefit of the multi-task set up, as a solution to missing data, is studied here in more depth by masking PPI interface labels for a part of the proteins in our dataset. We show that formulating a predication task as a multi-task learning problem can be beneficial for protein

structural prediction tasks for which only a small set of annotated training data is available.

Methods

The OPUS-TASS model described by Xu et al.¹⁶ is used as the basis for the prediction models used in this study. We also used their published annotated dataset.

Datasets

The combined OPUS-TASS training and validation set consists of 11,007 proteins and includes the following generated input features: HMM profiles, PSSM profiles, physicochemical features, and the PSP19 feature. The proteins in this dataset were selected by Hanson et al.⁴⁰, and were also used for training and validating SPOT-1D¹⁵. The proteins are culled from the PISCES⁴¹ server in February 2017. Only structures that are obtained by a resolution better than 2.5 Å and an R-free of less than 1.0 were selected. Sequences exceeding a sequence length of 700 residues were removed and the dataset was filtered by sequence identity, applying a cut-off of 25%. Residues based annotations for PPI interfaces are available for one third of this dataset (3,551 proteins). This PPI interface annotated data is a selection of PDB⁴² structures as described by Stringer et al.²². In short, the procedure was as follows. Proteins consisting of 2-200 chains were selected. For one structural complex the inter-atomic distance between all amino acids in separate chains were determined. The amino acids were defined as binding residues when the inter-atomic distance falls below a certain threshold. This threshold was set to 0.5 Å plus the Van der Waal's radii of the two amino acids.

Two datasets were constructed: 1) the 'PPI dataset' contains all proteins in the OPUS-TASS dataset for which PPI annotations were available; 2) the extended 'PPI_extendedSFD dataset' contains the PPI dataset extended with structural feature data stored in all other proteins of the OPUS-TASS dataset (see Supplementary Figure 1). The PPI_extendedSFD dataset is therefore larger, but only partially annotated with PPI interface information. Both datasets were split into a training (80%), validation (10%) and test (10%) set. For all these sets the PPI dataset is a strict subset of the PPI_extendedSFD dataset, and thus contains the same PPI interface information. Splitting the data into training, validation and test sets was performed *after* matching the proteins of OPUS-TASS dataset with PPI annotations based on the PDB ID and protein chain. The PPI annotations for 64 proteins had to be removed because the protein sequences between the two database did not correspond.

Input features

The sequence-based input features consist of 20 features obtained from the Position Specific Scoring Matrix (PSSM), 30 features obtained from Hidden Markov Model (HMM) profiles, 7 features obtained from the physicochemical properties and 19 features obtained from the PSP19 classification. Hence, every protein is represented by a matrix with the following dimensions: the number of protein residues times the 76 input features (20 + 30 + 7 + 19). The PSSM profiles, constructed by Xu et al.¹⁶, are based on three iterations of PSI-BLAST (v2.10.0+)⁴³ using the UniRef90 database⁴⁴. The HMM profiles are constructed using HHblits (v3.1.0)⁴⁵ and the Uniclust30 database⁴⁶. HMMs capture position-specific information about insertions and deletions, additional to conservation, for each amino acid. Only the conservation per amino acid is captured in the PSSM³¹. HHblits is a fast sequence search algorithm using HMM-HMM alignment after applying a profile-profile alignment as pre-filter⁴⁷. The 7 physicochemical properties are the amino acid properties as described by Meiler et al.⁴⁸. The PSP19 feature captures side-chain flexibility and packing orientation⁴⁹. For this feature 19 rigid-body blocks were constructed by Lu et al.⁴⁹ in a one-hot encoded sequence indicating the existence of a block in a protein residue. Both the physicochemical properties and the PSP19 feature are protein independent and amino acid specific.

Prediction tasks

During pre-processing, the prediction task labels were generated. For our models, we considered three output labels generated by Xu et al.¹⁶: secondary structure in 3 and 8 classes (S3 and S8), and the solvent accessibility of residues (SA). The three class secondary structure components are coil, α -helix and β -strand. These three components can be further distinguished into eight classes: coil into coil, high-curvature, and β -turn; α -helix into α -helix, 310-helix, and π -helix; and β -strand into β -strand and β -bridge⁵⁰. The S3, S8 and SA labels are derived from DSSP⁵⁰. In addition, we added two classification tasks: identification of buried (BU) and PPI interface (IF) residues. Residues were labelled as buried if the fraction of absolute solvent accessibility over its maximum solvent accessibility is less than 7%⁵¹. Residues with incomplete side chains were masked for the SA prediction. For the PPI labels all residues of a protein were masked when no PPI interface annotation was available for the entire protein. Masked residues were not taken into account in the loss calculation and performance measures. The prediction tasks S3, S8, SA, and BU are related to the PPI interface and therefore used in this study as the possible additional learning tasks for the multi-task IF predictor.

Model architecture

The details of the deep learning model architecture are described by Xu et al.¹⁶ in their Methods and Figure 1. The model is implemented in python using the keras library of tensorflow⁵². Similar to their model we performed data enhancement and used the architecture of 2 transformer layers, 5 CNN layers and 4 bidirectional LSTM layers (see Supplementary Figure 2). We used their dropout of 0.25 and the rectified linear unit (ReLU) activation function. In contrast to the OPUS-TASS model the possible output labels in our models are: interface residues (IF), secondary structure based on 3 classes (S3) and on 8 classes (S8), absolute solvent accessibility (SA), and buried residues (BU). The models studied are named based on the abbreviations of the prediction tasks that are considered in the model. We do not create an ensemble of these different models. Furthermore, we added the possibility to mask part of the PPI interface data, to study the effect of partially annotated data.

Multi-task learning process

The multi-task learning setup is implemented in the cost function of the model. For each prediction task the individual loss is determined by the cross-entropy for classification tasks (S3, S8, BU, IF) and mean squared error for the regression task (SA). The individual losses are summed and form the total cost function, as shown in Figure 2. In this way the model is able to learn a shared representation for all the prediction tasks. For the majority of the proteins in the PPI_extendedSFD dataset no PPI interface information is available. For these proteins the loss is constructed solely by the individual losses of the related tasks. We investigated the effect of the relative weights for each task in the cost function, using three approaches. Method A: the weights of the individual learning tasks losses in the total cost function were set equal. Method B: similar prediction tasks were grouped together (S3, S8 and BU, SA). The weights of pairs of similar prediction tasks were halved, such both sum up to a weight equal to the other tasks. Method C: a weight to the interface loss was assigned such that its fraction is 50% of the total cost function, i.e. the PPI task was given more weight in this approach.

Model settings

In line with Xu et al.¹⁶, the batch size was set to 4 proteins, initial weights were set by the glorot uniform initializer, and the Adam optimizer is used during training⁵³. The learning rate is divided by two when validation performance decreases, as measured by the area under the receiver operator characteristics curve (AUC ROC) of the PPI interface prediction.

We evaluated the early stopping criteria to avoid overfitting as defined for the OPUS-TASS model¹⁶ by training several models for 50 epochs. Xu et. al. defined the stopping criteria to converge when for the fourth time the AUC ROC score on the validation set is lower than the previous score. Training the

Figure 2. Visualisation of the possible protein structural prediction tasks, and the implementation of the multi-task setup in the cost function of the model. The protein in the example is Pterin-4- α -carbinolamine dehydratase 2, based on the PDB structure 4WIL chain A. **(a)** The protein-protein interaction interface residues are colored red for the protein chain (A), which is indicated in blue. These residues interact with the olive coloured protein chain. **(b)** Buried residues are indicated in green. **(c)** The absolute solvent accessibility is shown in blue colors. The darker the color of the residues the higher the solvent accessibility. **(d)** Classification of the secondary structure components in three classes. **(e)** Coil is colored green, α -helix yellow and β -strand red. Classification of the secondary structure components in eight classes. Coil is colored green, high-curvature blue, β -turn dark green, α -helix yellow, β -strand red and β -bridge dark red. Note that this protein does not have a 3_{10} -helix or a π -helix. The loss of the individual (possible) prediction tasks are summed in the cost function which is used during training of the multi-task model.

models further after reaching these stopping criteria did not show improvements in model performance. We therefore decided to use the same stopping criteria.

The single-task learner, called the IF model, is only trained on the PPI interface prediction task. This model was not able to identify interface residues when using an initial learning rate of $1e-3$, the initial learning rate used to train the OPUS-TASS model. After hyperparameter tuning, the initial learning rate was set to $2.5e-4$. Additionally, a weight was applied to the true interface residues in the loss calculation in order to adjust for the class imbalance of the PPI interface prediction. This weight was set to the ratio of non-interface residues over interface residues in the training and validation set. Therefore, this weight was set to 6.37.

Training and evaluation of the model is performed on one node containing a Titan X GPU. Models converge after approximately 10 epochs. On the PPI_extendedSFD dataset the duration of one epoch is approximately one hour. The validation performance of the model on all the prediction tasks together with training performances were collected with TensorBoard⁵².

Evaluation

We compare the performance of the single-task IF model to different multi-task models. The multi-task models contain different combinations of the related learning tasks as prediction tasks next to the interface prediction. After the models have seen all training sequences in one epoch, models are evaluated. We proceed training until the early stopping criteria, which is only based on the PPI interface prediction performance, is reached. Thereafter, the model outputs the highest reached AUC ROC score for the PPI interface prediction, as well as the corresponding scores of the related tasks, on the

validation set.

We use different performance measures for the different prediction tasks. The PPI interface prediction performance is determined by the AUC ROC. The secondary structure prediction in 3 classes and 8 classes, and the buried prediction performance is measured by the accuracy (ACC). The absolute solvent accessibility performance is measured by the Pearson correlation coefficient (PCC), which is a normalised measure of the covariance in the range between -1 and 1⁵⁴.

All models are trained four times, after which the mean performance and standard deviation is determined. The different models are compared based on the PPI AUC ROC scores. A one-sided test of significance is performed on the difference of the two independent AUC ROC scores⁵⁵, available from http://vassarstats.net/roc_comp.html.

We performed an error analysis on the individual proteins in the test set, to investigate the relation between small interfaces and IF performance scores. Furthermore, we test the relation between low IF prediction scores and the other structural feature task prediction scores, by performing a linear regression using the scipy.stats module (version: 1.3.1, see <https://docs.scipy.org/doc/scipy/reference/stats.html>).

Results

In order to test whether PPI interface prediction could be improved when formulated as a multi-task problem, we generated two datasets with structural and PPI annotations: 1) a PPI dataset that contains both structural and PPI annotations for all proteins and 2) a PPI_extendedSFD dataset that contains structural annotations for all proteins and PPI annotations for only a third of the data. Hence, the PPI dataset (3,551 proteins) is a subset of the PPI_extendedSFD dataset. The PPI_extendedSFD dataset (11,007 proteins) is augmented by the remaining protein structures in the OPUS-TASS dataset for which structural information – but not PPI annotations – are available, see also Table 1).

Structural information	PPI dataset		PPI_extendedSFD	
	SF + PPI	only SF	SF + PPI	only SF
Training set	2842	0	2842	5961
Validation set	353	0	353	749
Test set	356	0	356	746

Table 1. The number of proteins raining, validation and test sets for the PPI and PPI_extendedSFD datasets. The PPI dataset only contains proteins for which both structural features and PPI interface annotations are available. The PPI_extendedSFD dataset contains additional proteins for which only structural features are available.

Tuning the multi-task model

We adjusted the multi-task model setup of Xu et al.¹⁶ to allow masking of unannotated labels during the training process. We also included a performance measure for the absolute solvent accessibility prediction task. In addition, the model was extended to identify buried and PPI interface residues. These tasks were implemented in the same way as the existing secondary structure classification tasks.

The initial learning rate was tuned on the PPI dataset by training the IF model on the training set and validating the model on the validation set. We tuned this parameter by considering the values 1e-3, 5e-4, 2.5e-4, 1e-4, 7.5e-4, 5e-5 and 1e-5. Model performances was measured by AUC ROC, area under the precision-recall curve (AUC PR), and accuracy. A stable optimum is reached by training the model on the learning rate of 2.5e-4 (see Supplementary Figure 3). We evaluated the performance of the model on these learning rates, for the related tasks by training the best performing multi-task model

presented by Xu et al.¹⁶. The results show, in line with the PPI interface performance score, a stable optimum for a learning rate of 2.5×10^{-4} (see Supplementary Figure 4).

Performances of different multi-task learning models

We used the prediction tasks S3, S8, SA and BU as possible related learning tasks to the PPI interface prediction task. Several models were trained on different combinations of these tasks. Every model was trained four times on both the PPI and PPI_extendedSFD dataset separately after which the mean AUC ROC and AUC PR scores and their standard deviation was determined. Results, on the validation sets, are shown in Table 2. The multi-task learning strategy significantly ($P < 1 \times 10^{-3}$ for all models) outperforms the single task learner (AUC ROC: 73.17 ± 0.36) on both the PPI and the PPI_extendedSFD dataset. The 'IFBUS3SA' model, trained on the PPI_extendedSFD dataset using interface (IF), secondary structure in three classes (S3), buried (BU) and solvent accessibility (SA) as prediction tasks reaches the highest AUC ROC (76.32 ± 0.23). This multi-task model significantly outperforms the single-task model ($P < 1 \times 10^{-6}$) and the IFBU model on the PPI dataset ($P < 1 \times 10^{-3}$). Additionally, it outperforms the IFBU model on the PPI_extendedSFD dataset and the IFBUSA model on the PPI dataset significantly ($P < 0.01$). Including the more specific secondary structure classification task S8 instead of S3, or both S3 and S8, did not show further improvement.

	PPI dataset		PPI_extended dataset	
	AUC ROC	AUC PR	AUC ROC	AUC PR
IF	73.17 ± 0.36	31.71 ± 1.01	73.17 ± 0.36	31.71 ± 1.01
IFBU	74.85 ± 0.19	34.37 ± 0.32	75.15 ± 0.20	35.35 ± 0.17
IFBUSA	75.08 ± 0.24	35.62 ± 0.97	75.92 ± 0.21	36.65 ± 0.42
IFBUS3SA	75.73 ± 0.50	35.79 ± 1.44	76.32 ± 0.23	38.44 ± 0.92
IFBUS8SA	75.73 ± 0.31	36.39 ± 0.74	76.20 ± 0.24	37.95 ± 0.52
IFBUS3S8SA	75.73 ± 0.21	36.46 ± 1.13	76.06 ± 0.14	38.16 ± 0.93

Table 2. Comparison of the PPI interface performance of the single-task model against different multi-task models. The AUC ROC and AUC PR scores and the corresponding standard deviations are shown. Performance is measured on the validation set of both the PPI dataset and the augmented PPI_extendedSFD dataset. The multi-task models outperform the single task model (73.17 ± 0.36 AUC ROC) significantly on both dataset ($P < 0.001$). The overall highest AUC ROC score (76.32 ± 0.23), shown in bold, is reached when including buried residues, secondary structure in three classes and absolute solvent accessibility as related prediction tasks in addition to the PPI interface prediction on the PPI_extendedSFD dataset.

The PPI interface prediction performances expressed in AUC PR scores (see Table 2) follow similar trends to the AUC ROC scores. Further analysis showed that single-task learners of the related tasks optimised for that specific task reaches similar prediction performances as the multi-task learners optimised for PPI interface prediction (see Supplementary Figure 5). PPI interface prediction AUC ROC scores attained by additional models trained on more combinations of the related tasks can be found in Supplementary Figure 6.

We tried to improve the PPI interface predictions by including the torsion angle prediction as additional related learning task. We included the phi-angle and psi-angle prediction in all models described in Table 2. Performance was measured by the mean absolute error. However, after adding these tasks no significant improvement was shown compared to the best presented model (IFBUS3SA), see also Supplementary Table 1).

We tested our models on the independent test set. Results are shown in Figure 3 and illustrate the similar performances of the PPI interface prediction expressed in AUC ROC scores. These results

further support the conclusion that the multi-task learners outperform the single-task learner.

Figure 3. Comparison of the single-task model and multi-task models based on the AUC ROC scores of the PPI interface prediction on the validation and an independent test set. Performances are shown for the validation (dark blue) and test (blue) set for models trained on the PPI dataset, and the validation (dark red) and test (red) set for models trained on the PPI_extendedSFD dataset. All models are trained once on the training set. Similar performances are shown for the validation and test set. The multi-task models outperform the single-task model.

We investigated whether model performances could be improved by tuning the weights of the loss of the individual tasks in the total cost function using three different methods. Method A, for which the results are described above, weighs all tasks equally. Method B halves the weights of strongly related learning tasks (see Methods for further details). Method C keeps the contribution of the PPI interface prediction loss constant as 50% of the sum of the weights over all prediction tasks, this puts a much strong weight on the IF task. Both methods were tested during training of the IFBUS3SA, IFBUS8SA and IFBUS3S8SA models (see Supplementary Figure 7A). Method B was also compared to models including only one of the similar prediction tasks (see Supplementary Figure 7B). The results do not suggest an increase in model performance. Hence, the model does not appear to be very sensitive to the weights of the cost function.

Partially annotated dataset

When comparing the model performances with and without data extension (in Table 2 and Figure 3) we can observe a subtle increase in performance for the training datasets extended with structural property information, but without additional PPI interface annotations. These results suggest that the interface prediction benefits not only from the multi-task learning strategy by annotating the protein sequences in the PPI dataset by the related task information, but also from training on additional data of the related tasks only, as provided by the partially annotated dataset.

To further investigate these results we decreased the PPI interface annotations in the datasets. We evaluated the single-task learner and the best performing model IFBUS3SA on both datasets in which only a part of the data is considered, see Figure 4. We trained the single-task model and the IFBUS3SA model on a part of the PPI dataset. Next, we trained the IFBUS3SA model on the PPI_extendedSFD dataset for which we included only a part of the PPI annotations. At each data decreasing step, the three models are trained on the same PPI interface information and all models are evaluated on the total

validation set. Model performance is measured by AUC ROC for the PPI IF prediction.

Figure 4. The importance of the multi-task setup and the data extension when training a PPI interface prediction model on a small dataset. The single-task model IF (pink) and the multi-task model IFBUS3SA (red and brown) are compared. The IF model and the IFBUS3SA model indicated in red are trained on a part of the PPI dataset. Differences in performance between the pink and red bars therefore presents the benefit of the multi-task learning strategy. The IFBUS3SA model in brown is trained on the PPI_extendedSFD dataset in which only a part of the PPI interface information is considered. All the brown bars are thus trained on the same number of sequences for which the related task information is available. Differences in performance between red and brown bars indicate the benefit of training the model on the augmented PPI_extendedSFD dataset. Model performance is shown by the mean AUC ROC (bars) and standard deviation (whiskers) of the PPI interface prediction on the total validation set.

Figure 4 shows that, as expected, less training data generally leads to a worse performance for all strategies. For very small training datasets (e.g. one-twentieth of the total data) the multi-task learning (red bars), without extending the data, does not outperform the single task strategy (pink bars) significantly (Figure 4). This is probably also due to the lack of sufficient information to train on.

Figure 4 also shows that the smaller the training datasets the larger the difference in performance between the single-task strategy (pink bars) and multi-task strategy with data extension (brown bars). Hence, the performance of the multi-task model improves strongly when the training dataset is augmented with proteins only containing structural annotations, but for which the PPI annotations are missing, i.e. when it is trained on the PPI_extendedSFD dataset (brown compared to red). Thus, information captured by the related learning tasks improves PPI interface prediction even if PPI interface annotations are not available for the majority of the proteins in the training set. Likewise, the data extension becomes more important if annotated PPI interface data is very scarce. Equal performances are reached for the IF model trained on the entire available PPI interface information (pink bar labelled ‘all’) and the IFBUS3SA model trained on the PPI_extendedSFD dataset including only one eighth of the available PPI interface information (brown bar labelled ‘1/8’).

Error Analysis

We performed an error analysis on individual proteins in the test set after training the multi-task model IFBUS3SA to gain biological insight in the predictions generated by the models.

Figure 5 shows receiver operator characteristics curve for four single exemplary proteins. Target and predicted residues are shown using structure viewer UCSF ChimeraX⁵⁶ (see Figure 5).

The protein with high AUC ROC values show many correctly predicted interface residues (indicated in yellow in Figure 5 b, c and d). False positive residues, which are residues predicted as interface residues but not indicated as such in the gold standard (indicated in white in Figure 5b–e), are typically closely located to the actual interface for those (see Figure 5 b). Proteins corresponding to lower AUC ROC values show some false negatives, which are interface residues that are not predicted as interface residues by the model (indicated in red in Figure 5 c, d and e), and many false positives. Note that some false positives occur in localised regions of the structure, possibly indicating a true PPI interface region which is not annotated in the dataset. For example, in Figure 5 b the false positives located on the α -helices, actually form a secondary interface in the tetrameric structure. For one protein, with a very small interface, no correctly predicted residues were observed (see Figure 5 d).

Figure 5. Analysis of four proteins in the test set after training the multi-task model IFBUS3SA on the PPI_extendedSFD dataset. (a) The receiver operator characteristics (ROC) curve of the four proteins (4wilA in blue, 3rtlA in orange, 1vkca in green, and 4a0eA in red) with their corresponding area under the curve score (AUC ROC). (b, c, d and e) Visualised protein structures corresponding to the proteins in the ROC curve. The protein chains that contain the predicted interface are indicated in dark blue, the binding partners are indicated in green. Correctly predicted residues are colored yellow, false positive residues white and false negative residues red. (b) Protein structure of the protein 4wilA, corresponding to light blue line in the ROC curve. (c) Protein structure of the protein 3rtlA, corresponding to the orange line in the ROC curve. (d) Protein structure of the protein 1vkca, corresponding to the green line in the ROC curve. (e) Protein structure of the protein 4a0eA, corresponding to the red line in the ROC curve.

To analyse the error trends in more detail, we trained four different multi-task models, using different training rounds, and determined the mean performance per protein in the test set. Figure 6 A shows that proteins with low (<0.4) AUC ROC values for the PPI interface prediction are all proteins containing a small annotated interface region. However, proteins containing small interfaces do not necessarily result in low prediction scores. Moreover, we studied the relation between the PPI interface prediction and the related task predictions. This was done in order to test whether proteins corresponding to low AUC ROC scores are, in general, proteins for which structural features are hard to predict. The R^2 was determined and resulted in 0.010 for the buried residues, 0.016 for the secondary structure in three classes and 0.031 for the absolute solvent accessibility. Hence, no considerable correlation was found between the AUC ROC value of the PPI interface prediction and the related learning task for the IFBUS3SA model (see Figure 6B). The same conclusion was drawn after performing this analysis on

the IFBUS3SA model trained on only one tenth of the PPI interface information (See supplementary Figure 8).

Figure 6. Error analysis on the individual proteins in the test set after training the IFBUS3SA model on the PPI_extendedSFD dataset. **(a)** The mean AUC ROC of the interface prediction is plotted against the number of interface residues per protein. Low AUC ROC scores (<0.4) are only seen when the interface region of the protein is small (<20 residues). **(b)** The mean accuracy of the predicted structural features (BU, S3 and SA), and the mean Pearson correlation coefficient of the absolute solvent accessibility are plotted against the mean AUC ROC score of the PPI interface prediction per protein. Linear regression was performed resulting in an R^2 of 0.010 (BU), 0.016 (S3) and 0.031 (SA).

Discussion

Predicting the protein-protein interaction interface from sequence is a difficult task and annotations of interface residues are scarce. Here, we show how to overcome the problem of the limited size of the datasets by training a deep neural network predicting PPI interface residues using a multi-task learning strategy on a partially annotated dataset. All our multi-task models outperform the single-task model significantly ($P < 0.001$) on the PPI and the PPI_extendedSFD datasets. The single task model achieves AUC ROC: $73.2\% \pm 0.4$, while the best performing multi-task model reaches $76.3\% \pm 0.2$; this latter model includes as related prediction tasks the identification of buried residues, secondary structure and absolute solvent accessibility, in addition to predicting the PPI interface. Performances on the independent test set are in line with the results on the validation set. Hence the representations learned by the related structural annotation tasks can indeed help the model with the task of classifying PPI interface residues.

We show that the added benefit of the multi-task setup can be further increased by adding annotations for only the related tasks: the PPI interface prediction drastically improved when we extended the limited PPI training dataset with additional samples (proteins) for which only related structural annotations were available. Moreover, the multi-task setup becomes even more important when the training set is reduced. To highlight the strength of this, we removed all but one-eighth of the PPI interface information from the extended dataset. In this scenario the multi-task model still achieves similar performance scores as the single-task learner when trained on all PPI interface information. Hence, we show that formulating a predication task as a multi-task learning problem can be immensely powerful for protein structural (or functional) prediction tasks for which only a small set of annotated training data is available.

It is important to mention that in the current structural datasets, it is likely that many true protein-protein interaction (PPI) interface residues are not annotated as such, simply because no PDB structures of the relevant bound states are available. Hence, some true protein-protein binding interfaces will be

missing in any structural dataset used for training and performance evaluation of any PPI interface prediction method.

The scope of this study is to provide a proof of principle for the use of multi-task learning to improve prediction performance for protein structure tasks with scarce functional annotations, such as PPI interfaces. Our learning model is comparable to the multi-task model used in OPUS-TASS¹⁶. We did not perform extensive tuning of the architecture to obtain the highest possible PPI interface prediction accuracies. Nevertheless, our resulting AUC ROC scores are comparable to published state of the art methods for PPI interface prediction. Note that the performance of these different methods cannot be compared directly as different test sets were used. We will, however, include a discussion of the measured performance of these models to provide a background against which to better interpret our results, and to compare which features were included, and how these were used in the prediction model. A recent overview of state of the art PPI predictors is provided by Zhang & Kurgan⁹, which we recently extended with a comparison with some of our own methods²². These comparisons were all done on their ZK448 benchmark dataset¹³ using several metrics, below we list their AUC-ROC scores (unless noted otherwise). SSWRF by Wei et al.²⁹ achieved an AUC-ROC of 68.7%; it outperformed the state-of-art methods in 2016. Similarly to our multi-task model, this method used information of the PSSM and the solvent accessibility. In SSWRF these are both used as input feature, whereas in our model the latter is used as related learning task. Their third input feature, the averaged cumulative hydropathy, was not considered for our method. SeRenDIP by Hou et al.^{7,30} later achieved an AUC ROC score of 70.7% on a dataset of both homomeric and heteromeric protein complexes (HHC). SeRenDIP was not compared to SSWRF directly, but on the HHC test set did outperform older PPI prediction models SPPIDER⁵⁷ and PSIVER⁵⁸. SeRenDIP uses conservation, solvent accessibility, and secondary structure as input features to generate the PPI interface predictions. Furthermore, SeRenDIP includes backbone dynamics and sequence length, which are not used in our current multi-task model. SCRIBER achieves a AUC ROC of 71.5%, outperforming among others SPPIDER (51.7%), PSIVER (58.1%) and SSWRF on their ZK448 test set⁹. SCRIBER uses evolutionary conservation, relative solvent accessibility, and secondary structure features as input features. Furthermore, SCRIBER uses amino-acid physicochemical properties as input (charge, hydrophobicity, polarity, aliphaticity, aromaticity, acidity and size), as well as relative amino acid interface propensity and annotations of intrinsically disordered regions. Note that additional features, such as those used in other methods, may also be included in our multi-task model, either as input features or as related learning task. For example, we have previously shown that including sequence length and backbone flexibility as input features improves prediction of PPI⁷ and epitope regions⁵⁹. In addition, structure based features, such as (predicted) residue contacts, or other (predicted) 3D features could be added as input, or as related task. Recent advances in 3D structure prediction², suggests that this may be a likely way to further improve PPI interface prediction. However, the architecture of the model would have to be adapted. Several studies presented models to predict the interaction between proteins and other molecules like peptides, small molecules, and nucleic acids^{9,22,60}. These annotations could potentially be used as related tasks, and may further improve the model performance.

Although we did not perform extensive tuning of the model, we did investigate some parameters that are most likely to affect the multi-class learning. We implemented the multi-task learning strategy by the combined loss function which is used in all layers of the model (except the output layer). These layers are shared between all included tasks. We explored adjusting the weights of the different task, however, this did not result in any improvement of performance.

We furthermore tried to improve the PPI interface predictions by including torsion angle prediction, using the OPUS-TASS labels. However, no significant improvement was shown for the best presented model in this paper. In this work, we adjusted the initial learning rate, created our own additional output labels, and set weights for the combined loss function and the class imbalance for the PPI interface prediction. We hypothesize that similar PPI interface prediction performances could be achieved

when simplifying the model architecture. Further studies should conclude if the performance could be improved even further when fine-tuning the model. Nevertheless, here we show the substantial benefit of the multi-task learning strategy on a partially annotated dataset to achieving accurate performances for the difficult PPI interface prediction task. Therefore, we hope that the multi-task setup and the data extension will be of significant value in other protein structural or functional prediction tasks in which the size of annotated training sets is (extremely) limited.

References

1. Consortium, T. U. UniProt: the universal protein knowledgebase in 2021. *Nucleic Acids Res.* **49**, D480–D489, DOI: [10.1093/nar/gkaa1100](https://doi.org/10.1093/nar/gkaa1100) (2020).
2. Jumper, J. *et al.* Highly accurate protein structure prediction with AlphaFold. *Nature* **596**, 583–589 (2021).
3. Zhang, J. & Kurgan, L. Review and comparative assessment of sequence-based predictors of protein-binding residues. *Briefings Bioinforma.* **19**, 821–837, DOI: [10.1093/bib/bbx022](https://doi.org/10.1093/bib/bbx022) (2018).
4. Uetz, P. *et al.* A comprehensive analysis of protein–protein interactions in *saccharomyces cerevisiae*. *Nature* **403**, 623–627 (2000).
5. Jones, S. & Thornton, J. M. Principles of protein-protein interactions. *Proc. Natl. Acad. Sci.* **93**, 13–20 (1996).
6. Vinayagam, A. *et al.* A directed protein interaction network for investigating intracellular signal transduction. *Sci. signaling* **4**, rs8–rs8 (2011).
7. Hou, Q., Geest, P., Vranken, W. & Feenstra, K. A. Seeing the trees through the forest: Sequence-based homo-and heteromeric protein-protein interaction sites prediction using random forest. *Bioinformatics* **33**, 1479–1487, DOI: [10.1093/bioinformatics/btx005](https://doi.org/10.1093/bioinformatics/btx005) (2017).
8. Sanchez-Garcia, R., Sorzano, C. O. S., Carazo, J. M. & Segura, J. BIPSPi: a method for the prediction of partner-specific protein–protein interfaces. *Bioinformatics* **35**, 470–477 (2019).
9. Zhang, J. & Kurgan, L. SCRIBER: accurate and partner type-specific prediction of protein-binding residues from proteins sequences. *Bioinformatics* **35**, i343–i353 (2019).
10. Shoemaker, B. A. & Panchenko, A. R. Deciphering protein–protein interactions. part i. experimental techniques and databases. *PLOS Comput. Biol.* **3**, 1–8, DOI: [10.1371/journal.pcbi.0030042](https://doi.org/10.1371/journal.pcbi.0030042) (2007).
11. Kuzmanov, U. & Emili, A. Protein-protein interaction networks: probing disease mechanisms using model systems. *Genome medicine* **5**, 1–12 (2013).
12. Savojardo, C., Martelli, P. L. & Casadio, R. Protein–protein interaction methods and protein phase separation. *Annu. Rev. Biomed. Data Sci.* **3**, 89–112 (2020).
13. Zhang, Y. & Yang, Q. An overview of multi-task learning. *Natl. Sci. Rev.* **5**, 30–43, DOI: [10.1093/nsr/nwx105](https://doi.org/10.1093/nsr/nwx105) (2018).
14. Klausen, M. *et al.* NetSurfP-2.0: Improved prediction of protein structural features by integrated deep learning. *Proteins: Struct. Funct. Bioinforma.* **87**, 520–527, DOI: [10.1002/prot.25674](https://doi.org/10.1002/prot.25674) (2019).
15. Hanson, J., Paliwal, K., Litfin, T., Yang, Y. & Zhou, Y. Improving prediction of protein secondary structure, backbone angles, solvent accessibility and contact numbers by using predicted contact maps and an ensemble of recurrent and residual convolutional neural networks. *Bioinforma. (Oxford, England)* **35**, 2403–2410, DOI: [10.1093/bioinformatics/bty1006](https://doi.org/10.1093/bioinformatics/bty1006) (2019).

16. Xu, G., Wang, Q. & Ma, J. OPUS-TASS: a protein backbone torsion angles and secondary structure predictor based on ensemble neural networks. *Bioinformatics* **36**, 5021–5026, DOI: [10.1093/bioinformatics/btaa629](https://doi.org/10.1093/bioinformatics/btaa629) (2020).
17. Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. *Deep learning*, vol. 1 (MIT press Cambridge, 2016).
18. Heffernan, R. *et al.* Improving prediction of secondary structure, local backbone angles, and solvent accessible surface area of proteins by iterative deep learning. *Sci. Reports* **5**, 11476–11476, DOI: [10.1038/srep11476](https://doi.org/10.1038/srep11476) (2015).
19. Wang, S., Li, W., Liu, S. & Xu, J. RaptorX-property: A web server for protein structure property prediction. *Nucleic Acids Res.* **44**, gkw306, DOI: [10.1093/nar/gkw306](https://doi.org/10.1093/nar/gkw306) (2016).
20. Heffernan, R., Yang, Y., Paliwal, K. & Zhou, Y. Capturing non-local interactions by long short-term memory bidirectional recurrent neural networks for improving prediction of protein secondary structure, backbone angles, contact numbers and solvent accessibility. *Bioinformatics* **33**, 2842–2849, DOI: [10.1093/bioinformatics/btx218](https://doi.org/10.1093/bioinformatics/btx218) (2017).
21. Gao, Y., Wang, S., Deng, M. & Xu, J. Real-value and confidence prediction of protein backbone dihedral angles through a hybrid method of clustering and deep learning. *BMC Bioinforma.* **19**, DOI: [10.1186/s12859-018-2065-x](https://doi.org/10.1186/s12859-018-2065-x) (2018).
22. Stringer, B. *et al.* PIPENN: Protein interface prediction with an ensemble of neural nets. *bioRxiv* DOI: [10.1101/2021.09.03.458832](https://doi.org/10.1101/2021.09.03.458832) (2021).
23. Rao, R. *et al.* Evaluating protein transfer learning with TAPE. *Adv. neural information processing systems* **32**, 9689 (2019).
24. Madani, A. *et al.* Progen: Language modeling for protein generation. *arXiv preprint arXiv:2004.03497* (2020).
25. Caruana, R. Multitask learning. *Mach. learning* **28**, 41–75 (1997).
26. Pan, S. J. & Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge data engineering* **22**, 1345–1359 (2009).
27. Ofra, Y. & Rost, B. ISIS: interaction sites identified from sequence. *Bioinformatics* **23**, e13–e16 (2007).
28. Li, B.-Q., Feng, K.-Y., Chen, L., Huang, T. & Cai, Y.-D. Prediction of protein-protein interaction sites by random forest algorithm with mRMR and IFS. *PloS one* **7**, e43927 (2012).
29. Wei, Z.-S., Han, K., Yang, J.-Y., Shen, H.-B. & Yu, D.-J. Protein-protein interaction sites prediction by ensembling SVM and sample-weighted random forests. *Neurocomputing* **193**, 201–212 (2016).
30. Hou, Q. *et al.* SeRenDIP: SEquential REmasteriNg to DerIve profiles for fast and accurate predictions of PPI interface positions. *Bioinformatics* **35**, 4794–4796, DOI: [10.1093/bioinformatics/btz428](https://doi.org/10.1093/bioinformatics/btz428) (2019).
31. Zvelebil, M. J. & Baum, J. O. *Understanding bioinformatics* (Garland Science, 2007).
32. Ma, B., Elkayam, T., Wolfson, H. & Nussinov, R. Protein-protein interactions: Structurally conserved residues distinguish between binding sites and exposed protein surfaces. *Proc. Natl. Acad. Sci. United States Am.* **100**, 5772–5777, DOI: [10.1073/pnas.1030237100](https://doi.org/10.1073/pnas.1030237100) (2003).
33. Carl, N., Konc, J. & Janežic, D. Protein surface conservation in binding sites. *J. chemical information modeling* **48**, 1279–86, DOI: [10.1021/ci8000315](https://doi.org/10.1021/ci8000315) (2008).
34. Choi, Y. S., Yang, J.-S., Choi, Y., Ryu, S. H. & Kim, S. Evolutionary conservation in multiple faces of protein interaction. *Proteins* **77**, 14–25, DOI: [10.1002/prot.22410](https://doi.org/10.1002/prot.22410) (2009).

35. Chen, H. & Zhou, H.-X. Prediction of interface residues in protein–protein complexes by a consensus neural network method: test against NMR data. *Proteins: Struct. Funct. Bioinforma.* **61**, 21–35 (2005).
36. Hoskins, J., Lovell, S. & Blundell, T. L. An algorithm for predicting protein–protein interaction sites: abnormally exposed amino acid residues and secondary structure elements. *Protein Sci.* **15**, 1017–1029 (2006).
37. de Vries, S. J. & Bonvin, A. M. How proteins get in touch: interface prediction in the study of biomolecular complexes. *Curr. protein peptide science* **9**, 394–406 (2008).
38. Faber, H. & Matthews, B. A mutant t4 lysozyme displays five different crystal conformations. *Nature* **348**, 263–266 (1990).
39. Wright, P. E. & Dyson, H. J. Intrinsically unstructured proteins: re-assessing the protein structure-function paradigm. *J. molecular biology* **293**, 321–331 (1999).
40. Hanson, J., Paliwal, K., Litfin, T., Yang, Y. & Zhou, Y. Accurate prediction of protein contact maps by coupling residual two-dimensional bidirectional long short-term memory with convolutional neural networks. *Bioinformatics* **34**, 4039–4045 (2018).
41. Wang, G. & Dunbrack Jr, R. L. PISCES: a protein sequence culling server. *Bioinformatics* **19**, 1589–1591 (2003).
42. Berman, H. M. *et al.* The protein data bank. *Nucleic acids research* **28**, 235–242 (2000).
43. Altschul, S. F. *et al.* Gapped BLAST and PSI-BLAST: a new generation of protein database search programs. *Nucleic acids research* **25**, 3389–3402 (1997).
44. Suzek, B. E. *et al.* UniRef clusters: a comprehensive and scalable alternative for improving sequence similarity searches. *Bioinformatics* **31**, 926–932 (2015).
45. Steinegger, M. *et al.* HH-suite3 for fast remote homology detection and deep protein annotation. *BMC bioinformatics* **20**, 1–15 (2019).
46. Mirdita, M. *et al.* Uniclust databases of clustered and deeply annotated protein sequences and alignments. *Nucleic acids research* **45**, D170–D176 (2017).
47. Remmert, M., Biegert, A., Hauser, A. & Söding, J. HHblits: lightning-fast iterative protein sequence searching by HMM-HMM alignment. *Nat. methods* **9**, 173–175 (2012).
48. Meiler, J., Müller, M., Zeidler, A. & Schmäschke, F. Generation and evaluation of dimension-reduced amino acid parameter representations by artificial neural networks. *Mol. modeling annual* **7**, 360–369 (2001).
49. Lu, M., Dousis, A. D. & Ma, J. OPUS-PSP: an orientation-dependent statistical all-atom potential derived from side-chain packing. *J. molecular biology* **376**, 288–301 (2008).
50. Kabsch, W. & Sander, C. Dictionary of protein secondary structure: pattern recognition of hydrogen-bonded and geometrical features. *Biopolym. Orig. Res. on Biomol.* **22**, 2577–2637 (1983).
51. Hubbard, T. & Blundell, T. Comparison of solvent-inaccessible cores of homologous proteins: definitions useful for protein modelling. *Protein Eng. Des. Sel.* **1**, 159–171 (1987).
52. Abadi, M. *et al.* Tensorflow: A system for large-scale machine learning. In *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, 265–283 (2016).
53. Kingma, D. P. & Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

54. Benesty, J., Chen, J., Huang, Y. & Cohen, I. Pearson correlation coefficient. In *Noise reduction in speech processing*, 1–4 (Springer, 2009).
55. Hanley, J. A. & McNeil, B. J. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology* **143**, 29–36 (1982).
56. Pettersen, E. F. *et al.* UCSF ChimeraX: Structure visualization for researchers, educators, and developers. *Protein Sci.* **30**, 70–82 (2021).
57. Porollo, A. & Meller, J. Prediction-based fingerprints of protein–protein interactions. *Proteins: Struct. Funct. Bioinforma.* **66**, 630–645 (2007).
58. Murakami, Y. & Mizuguchi, K. Applying the naïve bayes classifier with kernel density estimation to the prediction of protein–protein interaction sites. *Bioinformatics* **26**, 1841–1848 (2010).
59. Hou, Q. *et al.* SeRenDIP-CE: sequence-based interface prediction for conformational epitopes. *Bioinformatics* DOI: [10.1093/bioinformatics/btab321](https://doi.org/10.1093/bioinformatics/btab321) (2021). Btab321.
60. Yan, J. & Kurgan, L. DRNApred, fast sequence-based method that accurately predicts and discriminates DNA- and RNA-binding residues. *Nucleic Acids Res.* **45**, e84–e84 (2017).
61. Bal, H. *et al.* A Medium-Scale Distributed System for Computer Science Research: Infrastructure for the Long Term. *Computer* **49**, 54–63, DOI: [10.1109/MC.2016.127](https://doi.org/10.1109/MC.2016.127) (2016).

Acknowledgements

We thank Peter Bloem for stimulating discussions and insightful advice on the learning setup. We kindly acknowledge usage of the Distributed ASCI Supercomputer DAS-5 at the VU University Amsterdam⁶¹.

Author contributions statement

S.A. designed the experiments. H.C. collected the datasets, implemented the methods and performed the experiments. H.C., K.A.F. and S.A. analysed and interpreted the results, wrote and revised the article text, and approved the final version for publication.

The authors declare that there is no conflict of interest.

Data Availability

The code is available at <https://github.com/ibivu/multi-task-PPI>. Note that all experimental protein structures used in this study have been deposited to the PDB⁴². The PDB accession codes, as well as the generated features and output labels used for training and validation is available at <https://ibi.vu.nl/downloads/multi-task-PPI/>.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [MultitaskPPIScientificReportssupplement.pdf](#)