



Published in final edited form as:

Proteins. 2021 June ; 89(6): 697–707. doi:10.1002/prot.26052.

Combination of deep neural network with attention mechanism enhances the explainability of protein contact prediction

Chen Chen, Tianqi Wu, Zhiye Guo, Jianlin Cheng*

Department of Electrical Engineering and Computer Science, University of Missouri, Columbia, MO 65211, USA

Abstract

Deep learning has emerged as a revolutionary technology for protein residue-residue contact prediction since the 2012 CASP10 competition. Considerable advancements in the predictive power of the deep learning-based contact predictions have been achieved since then. However, little effort has been put into interpreting the black-box deep learning methods. Algorithms that can interpret the relationship between predicted contact maps and the internal mechanism of the deep learning architectures are needed to explore the essential components of contact inference and improve their explainability. In this study, we present an attention-based convolutional neural network for protein contact prediction, which consists of two attention mechanism-based modules: sequence attention and regional attention. Our benchmark results on the CASP13 free-modeling (FM) targets demonstrate that the two attention modules added on top of existing typical deep learning models exhibit a complementary effect that contributes to prediction improvements. More importantly, the inclusion of the attention mechanism provides interpretable patterns that contain useful insights into the key fold-determining residues in proteins. We expect the attention-based model can provide a reliable and practically interpretable technique that helps break the current bottlenecks in explaining deep neural networks for contact prediction. The source code of our method is available at <https://github.com/jianlin-cheng/InterpretContactMap>

Keywords

protein structure prediction; predicted contact map; neural network

1. Introduction

Prediction of residue-residue contacts in proteins plays a vital role in the computational reconstruction of protein tertiary structure. Recently, advancements in the mathematical and statistical techniques for inter-residue coevolutionary analysis provide essential insights for correlated mutation-based contact prediction, which is now becoming a critical component to generate input features for machine learning contact prediction algorithms. For instance, in the recent 13th Community-Wide Experiment on the Critical Assessment of Techniques for Protein Structure Prediction (CASP13) contact prediction challenge, significant

*corresponding author: Contact information for corresponding author: Office: 109 Naka Hall, Lab: 304, 305, 307 Engineering Building North, 250 Naka Hall, Phone: 573-882-7306, Fax: 573-882-8318, chengji@missouri.edu.

improvements have been achieved due to the integration of both inter-residue coevolutionary analysis and novel deep learning architectures¹⁻⁴.

A variety of deep learning-based models have been proposed to improve the accuracy of protein contact prediction since deep learning was applied to the problem in 2012 CASP10 experiment⁵. Many of these methods leverage the contact signals derived from the direct coupling analysis (DCA). Most DCA algorithms⁶⁻⁹ generate correlated mutation information between residues from multiple sequence alignments (MSAs), which is utilized by the deep convolutional neural networks in the format of 2D input feature maps. For example, RaptorX-Contact¹⁰, DNCON2¹¹, and MetaPSICOV¹² are a few early methods that apply the deep neural network architectures with one or more DCA approaches for contact prediction. The connection between the two techniques underscores the importance of explaining the contribution of patterns in coevolutionary-based features to the deep learning-based predictors.

Despite the great success of deep learning-based models in a variety of tasks, this approach is often treated as black-box function approximators that generate classification results from input features. Since the number of parameters in a network is somewhat proportional to its depth, it is infeasible to extract human-understandable justifications from the inner mechanisms of deep learning without proper strategies. Saliency maps and feature importance scores are widely used approaches for model interpretation in machine learning. However, due to the unique characteristic of contact prediction, these methods involve additional analysis procedures that require far more computational resources than other typical applications. For example, the saliency map for a protein with length L requires $L \times L$ times of deconvolution operations in a trained convolutional neural network since the output dimension of contact prediction is always the same as its input. While this number can be reduced by choosing only positive labels for analysis, the whole saliency map is still much harder to determine since the many DCA features fed to the network have higher dimensions than the traditional image data. For example, RaptorX-Contact¹⁰, one of the state-of-the-art contact predictors, takes 2D input with a size of $L \times L \times 153$. The recent contact/distance predictor DeepDist¹³ takes input with size up to $L \times L \times 484$. The very large input size for contact prediction makes it difficult to use these model interpretation techniques.

Recently, the attention mechanism has been applied in natural language processing (NLP)^{14,15}, image recognition¹⁶, and bioinformatics^{17,18}. The attention mechanism assigns different importance scores to individual positions in its input or intermediate layer so that the model can focus on the most relevant information anywhere within the input. In 2D image analysis, the attention weights for any individual positions on an image allow the visualization of critical regions that contribute to the final predictions.

In addition, these weights are generated during the inference step, without requiring additional computation procedures after the prediction of a contact map. Hence, the attention mechanism is a suitable technique to facilitate the interpretation of protein contact prediction models.

In this article, we propose an attention-equipped deep learning method for protein contact prediction, which adopts two different architectures of the attention targeted for interpreting 2D and 1D input features, respectively. The regional attention utilizes the $n \times n$ region around each position of its input 2D map while the sequence attention utilizes the whole range of its 1D input. The regional attention module is implemented with a specially designed 3D convolutional layer so that training and prediction on large datasets can be performed with high efficiency. The sequence attention is similar to the multi-headed attention mechanism applied in the NLP tasks. We show that by applying attention mechanisms on the general deep learning predictors, we can acquire models that are able to explain how position-wise information anywhere in input or hidden features are transferred to later contact predictions, and this can be done without significant extra computational cost and decrease of the prediction accuracy.

2 Materials and Methods

2.1 Overview

The overall workflow of this study is shown in Figure 1. We use the combined predictions from two neural network modules of different attention mechanisms (sequence attention and regional attention) to predict the contact map for a protein target. Both modules take two types of features as inputs: the pseudolikelihood maximization matrix (PLM) ⁸ generated from multiple sequence alignment as a coevolution-based 2D feature and the position-specific scoring matrix (PSSM) which provides the representation of the sequence profile for the input protein sequence. The outputs of the two modules are both $L \times L$ contact maps with scores ranging from 0 to 1, where L represents the length of the target protein. The final prediction is produced by the ensemble of two attention modules. We implemented our model with Keras (<https://keras.io>). For the evaluation of the predicted contacted contact map, we primarily focus on long-range contacts (sequence separation between two residues: $n - 24$).

2.2 Datasets

We select targets from the training protein list used in DMPfold ¹⁹ and extract their true structures from the Protein Data Bank (PDB) to create a training dataset. After removing the redundant proteins that may have >25% sequence identity with any protein in the validation dataset and test dataset, 6463 targets are left in the training dataset. The validation set contains 144 proteins used to validate DNCON2 ¹¹. The blind test dataset is built from 31 CASP13 free modeling (FM) domains. The CASP13 test dataset contains new proteins that have no sequence similarity with both the training and test datasets at all.

2.3 Input Feature Generation

For each protein sequence, we use two features as inputs for the deep learning model: PLM and PSSM. The PLM is generated from multiple sequence alignments (MSAs) produced by DeepMSA [16]. The sequence databases used in the DeepMSA homologous sequences search include Uniclust30 (2017–10) ²⁰, Uniref90 (2018–04) and Metaclust50 (2018–01) ²¹, our in-house customized database which combines Uniref100 (2018–04) and metagenomics sequence databases (2018–04), and NR90 database (2016). All of the sequence databases

used for feature generation were constructed before the CASP13 experiment (e.g. before the CASP13 test dataset was created). DeepMSA combines iteratively homologous sequence search of HHblits²² and Jackhmmer²³ on the sequence databases to compute MSAs for feature generation. It performs trimming on the sequence hits from Jackhmmer with a sequence clustering strategy, which reduces the search time of the HHblits database construction for the next round of search. The final input of the model consists of two major components: 1D features (PSSM) of dimension $L \times 20$ and 2D features (PLM) of dimension $L \times L \times 441$.

2.4 Deep Network Architectures

Our model consists of two primary components, the regional attention module, and the sequence attention module (Figure 1). The two modules include the attention layers, normalization layers, convolution layers and residual blocks. The outputs of these two modules are combined to generate the final prediction. Below are the detailed descriptions of each module with an emphasis on the attention layers.

2.4.1 Sequence attention module—In the sequence attention module (Figure 1), the 1D PSSM feature first goes through an instance normalization layer²⁴ and a 1D convolution operation, which is followed by a Bi-Directional long- and short-term memory network (LSTM) in which the LSTM operations are applied on both forward and reverse directions of the inputs. The output vectors on both directions are concatenated. The outputs are then fed into a multi-headed scaled dot product attention layer (Figure 2a). Three vectors required for the attention mechanism: Q(Query), K(Key), and V(Value) are generated from different linear transformations of the input of the attention layer. The attention output Z is computed as:

$$Z = \text{Softmax}\left(\frac{Q \times K^T}{\sqrt{d_{att}}}\right) \times V$$

where d_{att} represents the dimensions of Q, K and V. The attention 1D operation assigns different weights to the transformed 1D feature so that the critical input region for the successful prediction can be spotted. The attention output Z is then tiled to 2D format by the repetition of columns and rows for each dimension.

The 2D feature PLM first goes into the instance normalization and a ReLU activation²⁵. It is then processed by a convolutional layer with 128 kernels of size 1×1 and a Maxout layer²⁶ to reduce the input dimension from 128 to 64. The 2D inputs are concatenated with the tiled attention output and go into the residual network component. The final output of the sequence attention module is generated from a 2D convolution layer with a filter of size (1,1) and Sigmoid activation, resulting in output of size $L \times L$.

2.4.2 Regional attention module—The regional attention module (Figure 1) takes inputs from the PLM matrix and the tiled 2D PSSM feature. The two features are concatenated at the beginning of the module and are processed in the same way as the 2D PLM input of the sequence attention module. The residual network component with the

same configuration (described in 2.4.3) as in the sequence attention module is also applied. The last residual block is followed by a convolutional layer with 32 filters, and the results are used as the input of the attention 2D layer.

The input shape of the attention 2D layer (Figure 2b) is $(L, L, 32)$. It is converted by a 3D convolution layer (Region Stretching layer) with specially designed filters so that the output has shape $(L, L, 32, n^2)$, where n is the dimension of the attention region for each position in the 2D input. The purpose of this layer is to make the last dimension of its output represent the flattened n by n region around each element of the original input (in our model n is set to 5). Thus, each position in the L by L output are determined by the weighted sum of the n by n square window around itself. The Region Stretching layer has n^2 filters with shape $n \times n$. For the i -th filter of the layer, the weight of the i -th element (flatten in row-major order) in the $n \times n$ area is always set to 1 with all other positions set to 0. We repeat these filters 32 times so that the stretching operation is applied to all dimensions of the input. The weights of these filters will not be changed during training. This operation can leverage the highly optimized convolution implementation in Keras (https://keras.io/api/layers/convolution_layers/convolution3d/) and is much more efficient than the explicit implementation. The corresponding Q, K, and V vectors for the attention mechanism are computed from the transformed output of 3D convolution. The scaling and Softmax normalization are applied to the last dimension for the products of Q and K so that different attention weights can be assigned to the $n \times n$ surrounding area for each position on the $L \times L$ map. As a result, the output of each position on the feature map will be a weighted sum of their surrounding regions. After the 2D attention layer, the output of the regional attention module is generated from a 2D convolutional layer with a filter of size (1,1) and the Sigmoid activation.

2.4.3 Residual network architecture—Both attention modules have the same residual network component consisting of four residual blocks differing in the number of internal layers (Figure 1). Each residual block is composed of several consecutive instance normalization layers and convolutional layers with 64 kernels of size 3×3 . The number of layers showed in each block represents the number of 2D convolution layers in the corresponding component. The final values of the last convolutional layer are added to the output of a shortcut block, which is a convolutional layer with 64 kernels of size 1×1 . A squeeze-and-excitation (SE) block²⁷ is added at the end of each residual block. The SE operation weights each of its channels differently by a trainable 2-layer dense network when creating the output feature maps, so that channel-wise feature responses can be adaptively recalibrated.

2.5 Training

The training of the deep network is performed with the customized Keras data generators to reduce the memory requirement. The batch size is set to 1 due to the large size of feature data produced from long protein sequences. A normal initializer²⁸ is used to initialize the weights of the layers in the network. Adam optimizer²⁹ is used for training, with the initial learning rate set to 0.001. For epochs (complete passes through the entire training data) 30, the optimizer is switched to stochastic gradient descent (SGD), with learning rate and

momentum set to 0.01 and 0.9, respectively. The learning rate determines the scale for model parameters update at each iteration and the momentum³⁰ is used to compute the next update of the weights as a linear combination of the current gradient and the update of corresponding weights in the previous iteration. At the end of each epoch, the current weights are saved, and the precision of top L/2 long-range contact predictions (e.g. predicted contacts with sequence separation ≥ 24) on the validation dataset is evaluated. The training process is terminated at epoch 60, and the epoch with the best performance on the validation dataset is chosen for the final blind test.

3 Results

3.1 The contact prediction accuracy on the CASP13 dataset

We evaluate our models on 31 CASP13 free-modeling (FM) targets based on the average of the per-target performance on them. According to the definition from CASP13, a pair of residues are considered to be in contact if the distance between their C_{β} atoms in the native structure is less than 8.0 Å. By convention, long-range contacts are defined as contact pairs in which the sequence separations between the two residues of the contacts are larger than or equal to 24 residues. The sequence separation for medium-range is between 12 and 23 and short-range between 6 and 11 residues. Following a common standard in the field¹, we evaluate the precision of top L/n ($n = 1, 2, 5$) predicted long-range contacts. In addition to evaluating the overall performance of the combined model, we benchmarked the predictions from the two independent attention modules. The evaluation results are shown in Table 1.

The combined model outperforms each individual attention model and model without attention mechanism for top L/5, top L/2, and top L predicted contacts in medium and long-range. For instance, the top L/5 long-range precision of the combined model is 70.73%, higher than both the sequence attention module (64.46%) and the regional attention module (67.32%) as well as the baseline model that without either of the attention mechanisms. According to the pair t-test, the combined model performance is significantly better than the sequence model in all ranges ($p < 0.05$), while no significant difference is observed when compared with the baseline or regional attention model. We also compare the performance of our method with the top 10 methods in CASP13 on the FM targets (Table 2) and show that it achieves the overall performance comparable to the top-ranked CASP13 methods. Specifically, the accuracy of top L/5 or top L/2 predictions of our method (“Combined Attention”) is ranked 2nd out of the 11 methods.

We also find that the predictive improvements in combining the two attention modules are from the predictions with high confidence scores. Figure 3a and 3b illustrate the receiver operating curve (ROC) and Precision-Recall curves (PR curve) of the three models on targets for evaluation. The area under the curve (AUC) for ROC curve and PR curve of all three models has similar trends. Figure 3c and 3d show the ROC and PR curves of the union of residue pairs from top-L/5 scores in any of the three models. For AUC of both curves, the combined results have a higher score (0.7888 for ROC curve and 0.8031 for PR curve) than the sequence attention model (0.7614 for ROC curve and 0.7935 for PR curve) and the regional attention model (0.7769 for ROC curve and 0.7907 for PR curve). The improved

performance of combining the two attention models indicates that the ensemble of two different attention architectures enhances the final prediction.

For both models, we have also benchmarked the impact of different combinations of attention configuration (number of attention heads and size of attention regions) with the maximum scale of the architectures allowed by our GPU memory capacity (Nvidia GeForce 1080Ti 11G). The results are showed in Table 3.

3.2 Comparison of the predictive performance of two attention modules

We compare the performance of the two attention modules for each target in Figure 4. The results show that the precision scores of the two attention modules have a strong correlation (Pearson Correlation Coefficient=0.78) among all targets. As expected, most of the targets with high prediction precision in the combined model are those with high precision scores in both attention modules. Interestingly, there are cases in which the combined predictions acquire an improved performance when the two attention modules perform very differently. For example, the top-L/5 precision score of T1008_D1 reaches 93.33% in the combined model, higher than the sequence module (46.67%) and the regional module (80.00%). Similarly, the top-L/5 precision score of T0957s2 reaches 64.52% in the combined model, which is equal to the sequence module and higher the regional module (45.16%). These results further confirm that the difference in the architecture of two attention mechanisms provides a complementary effect that can contribute to performance improvement.

3.3 Visualization of attention scores from the sequence model

Our sequence attention model is similar to the neural translation model proposed in Google Transformer¹⁴, in which the attention weights are visualized through case studies of the importance of each word in the source language for a sentence to each word in the target language. While it may be infeasible to directly understand the importance of each residue in a protein in the folding process through human observations, we included several proteins (2PTL, 1IDY and 1SHG)³¹⁻³⁴ that have been studied through experimentally determined Φ -values. The Φ -values are the ratio of the change in stability of the transition-state ensemble (TSE) to that of the native state during folding due to the mutation of each residue, and represent important information about residue interactions present within the TSE³⁵.

Next, we demonstrate how position-wise information in sequence attention model is transferred to later-stage contact predictions in our attention mechanisms for 1D features. Since the sequence attention score is a L by L matrix, in which element (i,j) represents the importance of the j -th residue to the i -th residue, and the sum of each row is normalized to 1. Thus, the column sum of the attention weights can represent the overall importance of each residue according to the 1D input. We first checked the column sums of attention scores of these three proteins and compared the density of scores from regions of the highest Φ -value peak and scores from the rest regions. The results are shown in Figure 5. We show that the scores of Φ -value peak are significantly higher than other regions (p-value < 0.01, Wilcoxon test).

We further show the case studies of the patterns of attention scores for each protein (Figure 6). For each of the protein, we could identify at least one attention head (out of 4) that have a

clear pattern of the Φ -value peak residue from the experiment. In addition, this pattern is generated from the 1D feature of PSSM only, while in PSSM such a pattern is not observed at all.

3.4 Regional Attention scores and key residue pairs in successful prediction

We first consider the importance of the area with the high attention scores in contact prediction. To demonstrate this, we permute the input features around the positions that have high or low attention scores and use this permuted feature for prediction. Our results show that the number of true positive predictions will decrease most drastically (Figure 7), indicating that they contain important information related to protein fold. Also, the level of decrease remains similar when the region of permuted data grows from 1×1 to 5×5 in areas with high attention scores. In contrast, the level of decrease in areas with low attention score is much smaller and increases with the expansion of the permuted area. These results indicate the existence of potential protein folding-related key information in small areas with high attention scores.

To further explore the interpretability of our method, we analyze the model on a protein whose folding mechanism has been well studied: Human common-type acylphosphatase (AcP). The structure and sequence information of AcP is obtained from PDB (<https://www.rcsb.org/structure/2W4C>). Vendruscolo et al.³⁶ identified three key residues in AcP (Y11, P54, and F94) that can form a critical contact network and result in the folding of a polypeptide chain to its unique native-state structure. The 3D structure model and three key residues are shown in Figure 8a.

We use the regional attention module to predict the contact map of the protein. The precisions of the top-L/5, L/2, and L prediction are 100%, 95.74%, and 75.79%, respectively. We then extract the 2D attention score matrix from the model and combine the normalized row sums and column sums to reformat its dimension to $L \times 1$. The attention score mapped to the protein 3D structure spot two key residues: Y11 and F94, where large regions of high attention weights are located (Figure 8b). Furthermore, we apply the same strategy with the experimentally determined Φ -values on the 3D structure of AcP (Figure 8c). The comparison (Figure 8d, 8e) shows that the Φ -values and normalized attention scores have similar trends along the peptide sequence (Pearson correlation coefficient = 0.4) with three peaks for Y11, P54, and F94 appeared in neighboring regions of the curves determined by both the experimental method and the attention method. Also, we find that the true contact map does not provide the same level of information about the three key residues (Figure 8f). These results indicate that the attention scores can be applied to identify the critical components of the input feature. However, we also find that the co-evolutionary input scores calculated by PSICOV can also be used to identify some Φ -value peaks of AcP. Therefore, the 2D regional attention weights can be either a new way to identify folding-related new residues or summarization of the input. This situation is different from the 1D sequence attention, where the 1D attention weights can definitely identify Φ -value peaks (folding-related residues) that cannot be recognized from 1D inputs at all. Therefore, attention mechanisms can improve the explainability of contact prediction models, but the effects are not guaranteed and may depend on their architecture and inputs.

4 Discussion

Attention mechanisms have two valuable properties that are useful for protein structure prediction. First, attention mechanisms can identify important input or hidden features that are important for structure prediction, and therefore they have the potentials to explain how predictions are made and even increase our understanding of how proteins may be folded. However, the knowledge gained from the attention mechanisms depends on how they are designed and the input information used with them. Second, attention mechanisms can pick up useful signals relevant to protein structure (e.g. contact) prediction anywhere in the input, which is much more flexible than other deep neural network architectures such as sequential information propagation in recurrent neural networks and spatial information propagation in convolutional neural networks. As protein folding depends on residue-residue interactions that may occur anywhere in a protein, the attention mechanisms can be a natural tool to recognize the interaction patterns relevant to protein structure prediction or folding more effectively.

5 Conclusion

Interrogating the input-output relationships for complex deep neural networks is an important task in machine learning. It is usually infeasible to interpret the weights of a deep neural network directly due to their redundancy and complex nonlinear relationships encoded in the intermediate layers. In this study, we show how to use attention mechanisms to improve the interpretability of deep learning contact prediction models without compromising prediction accuracy. More interestingly, patterns relevant to key fold-determining residues can be extracted with the attention scores. These results suggest that the integration of attention mechanisms with existing deep learning contact predictors can provide a reliable and interpretable tool that can potentially bring more insights into the understanding of contact prediction and protein folding.

Acknowledgments

Funding

This work was supported by Department of Energy grants (DE-SC0020400 and DE-SC0021303), NSF grants (DBI1759934 and IIS1763246), and NIH grants (GM093123).

Reference

1. Shrestha R, Fajardo E, Gil N, et al. Assessing the accuracy of contact predictions in CASP13. *Proteins*. 2019;87(12):1058–1068. [PubMed: 31587357]
2. Senior AW, Evans R, Jumper J, et al. Improved protein structure prediction using potentials from deep learning. *Nature*. 2020;577(7792):706–710. [PubMed: 31942072]
3. Hou J, Wu T, Cao R, Cheng J. Protein tertiary structure modeling driven by deep learning and contact distance prediction in CASP13. *Proteins*. 2019;87(12):1165–1178. [PubMed: 30985027]
4. Zheng W, Li Y, Zhang C, Pearce R, Mortuza SM, Zhang Y. Deep-learning contact-map guided protein structure prediction in CASP13. *Proteins*. 2019;87(12):1149–1164. [PubMed: 31365149]
5. Eickholt J, Cheng J. Predicting protein residue-residue contacts using deep networks and boosting. *Bioinformatics*. 2012;28(23):3066–3072. [PubMed: 23047561]

6. Morcos F, Pagnani A, Lunt B, et al. Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proceedings of the National Academy of Sciences*. 2011;108(49):E1293–E1301.
7. Ekeberg M, Lovkvist C, Lan Y, Weigt M, Aurell E. Improved contact prediction in proteins: using pseudolikelihoods to infer Potts models. *Phys Rev E Stat Nonlin Soft Matter Phys* 2013;87(1):012707. [PubMed: 23410359]
8. Seemayer S, Gruber M, Soding J. CCMpred--fast and precise prediction of protein residue-residue contacts from correlated mutations. *Bioinformatics*. 2014;30(21):3128–3130. [PubMed: 25064567]
9. Jones DT, Buchan DW, Cozzetto D, Pontil M. PSICOV: precise structural contact prediction using sparse inverse covariance estimation on large multiple sequence alignments. *Bioinformatics*. 2012;28(2):184–190. [PubMed: 22101153]
10. Wang S, Sun S, Xu J. Analysis of deep learning methods for blind protein contact prediction in CASP12. *Proteins*. 2018;86 Suppl 1:67–77. [PubMed: 28845538]
11. Adhikari B, Hou J, Cheng J. DNCON2: improved protein contact prediction using two-level deep convolutional neural networks. *Bioinformatics*. 2018;34(9):1466–1472. [PubMed: 29228185]
12. Agranoff D, Fernandez-Reyes D, Papadopoulos MC, et al. Identification of diagnostic markers for tuberculosis by proteomic fingerprinting of serum. *The Lancet*. 2006;368(9540):1012–1021.
13. Wu T, Guo Z, Hou J, Cheng J. DeepDist: real-value inter-residue distance prediction with deep residual convolutional network. *bioRxiv*. 2020:2020.2003.2017.995910.
14. Vaswani A, Shazeer N, Parmar N, et al. Attention is all you need. Paper presented at: *Advances in Neural Information Processing Systems*2017.
15. Yang Z, Yang D, Dyer C, He X, Smola A, Hovy E. Hierarchical attention networks for document classification. Paper presented at: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*2016.
16. Xu K, Ba J, Kiros R, et al. Show, attend and tell: Neural image caption generation with visual attention. Paper presented at: *International conference on machine learning*2015.
17. Hu Y, Wang Z, Hu H, et al. ACME: pan-specific peptide-MHC class I binding prediction through attention-based deep neural networks. *Bioinformatics*. 2019;35(23):4946–4954. [PubMed: 31120490]
18. Chen C, Hou J, Shi X, Yang H, Birchler JA, Cheng J. Interpretable attention model in transcription factor binding site prediction with deep neural networks. 2019:648691.
19. Greener JG, Kandathil SM, Jones DT. Deep learning extends de novo protein modelling coverage of genomes using iteratively predicted structural constraints. *Nat Commun* 2019;10(1):3977. [PubMed: 31484923]
20. Mirdita M, von den Driesch L, Galiez C, Martin MJ, Soding J, Steinegger M. Uniclust databases of clustered and deeply annotated protein sequences and alignments. *Nucleic Acids Res* 2017;45(D1):D170–D176. [PubMed: 27899574]
21. Steinegger M, Soding J. Clustering huge protein sequence sets in linear time. *Nat Commun* 2018;9(1):2542. [PubMed: 29959318]
22. Remmert M, Biegert A, Hauser A, Soding J. HHblits: lightning-fast iterative protein sequence searching by HMM-HMM alignment. *Nat Methods*. 2011;9(2):173–175. [PubMed: 22198341]
23. Mistry J, Finn RD, Eddy SR, Bateman A, Punta M. Challenges in homology search: HMMER3 and convergent evolution of coiled-coil regions. *Nucleic Acids Res* 2013;41(12):e121. [PubMed: 23598997]
24. Ulyanov D, Vedaldi A, Lempitsky V. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:160708022*. 2016.
25. Nair V, Hinton GE. Rectified linear units improve restricted boltzmann machines. Paper presented at: *Proceedings of the 27th international conference on machine learning (ICML-10)*2010.
26. Goodfellow IJ, Warde-Farley D, Mirza M, Courville A, Bengio Y. Maxout networks. *arXiv preprint arXiv:13024389*. 2013.
27. Hu J, Shen L, Sun G. Squeeze-and-excitation networks. Paper presented at: *Proceedings of the IEEE conference on computer vision and pattern recognition*2018.

28. He K, Zhang X, Ren S, Sun J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. Paper presented at: Proceedings of the IEEE international conference on computer vision2015.
29. Kingma DP, Ba J. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980. 2014.
30. Qian N. On the momentum term in gradient descent learning algorithms. *Neural Networks*. 1999;12(1):145–151. [PubMed: 12662723]
31. Kim DE, Fisher C, Baker D: A breakdown of symmetry in the folding transition state of protein L11 Edited by C. R. Matthews. *Journal of Molecular Biology* 2000, 298(5):971–984. [PubMed: 10801362]
32. Gianni S, Guydosh NR, Khan F, Caldas TD, Mayor U, White GW, DeMarco ML, Daggett V, Fersht AR: Unifying features in protein-folding mechanisms. *Proceedings of the National Academy of Sciences of the United States of America* 2003, 100(23):13286–13291. [PubMed: 14595026]
33. Martínez JC, Serrano L: The folding transition state between SH3 domains is conformationally restricted and evolutionarily conserved. *Nature Structural Biology* 1999, 6(11):1010–1016. [PubMed: 10542091]
34. Wako H, Abe H: Characterization of protein folding by a Φ -value calculation with a statistical-mechanical model. *Biophysics and physcobiology* 2016, 13:263–279. [PubMed: 28409079]
35. Lindorff-Larsen K, Paci E, Serrano L, Dobson CM, Vendruscolo M. Calculation of mutational free energy changes in transition states for protein folding. *Biophys J* 2003;85(2):1207–1214. [PubMed: 12885664]
36. Vendruscolo M, Paci E, Dobson CM, Karplus M. Three key residues form a critical contact network in a protein folding transition state. *Nature*. 2001;409(6820):641–645. [PubMed: 11214326]

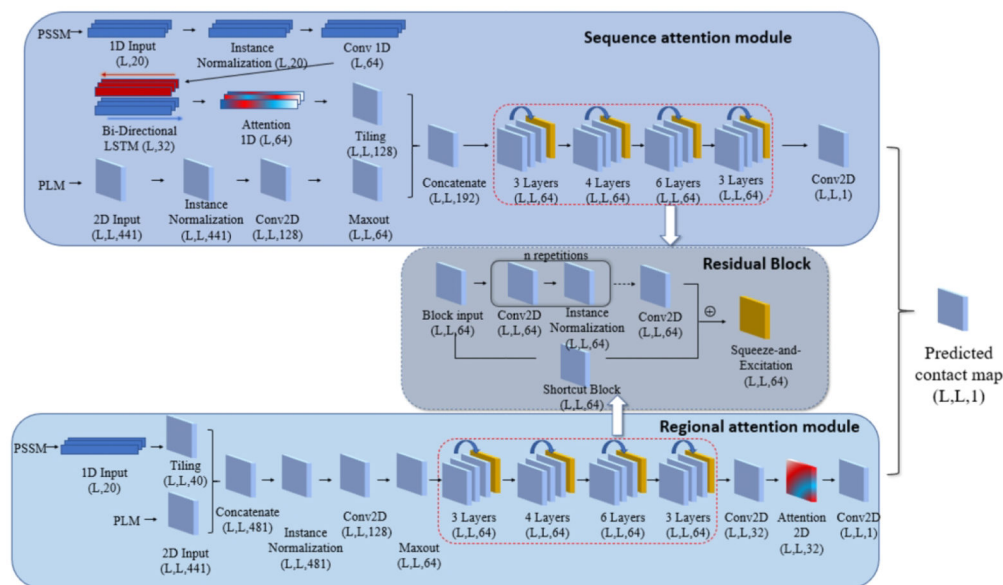


Figure 1.

An overview of the proposed attention mechanism protein contact predictor framework. The architecture of the deep neural network employed with two attention modules: In the sequence attention module, the 1D input (PSSM) first goes through the 1D convolution network and bidirectional LSTM. Then the attention mechanism is applied to the LSTM output. The 2D input (PLM) is first processed with the 2D convolutional neural network and the Maxout layer. The 1D input is then tiled to 2D format so that it can be combined with the 2D input. The concatenated inputs then go through a residual network with 4 residual blocks consist of 3, 4, 6, 3 repetitions of 2D convolution layers, respectively. In regional attention networks, the 1D inputs are first tiled to 2D format and concatenated with the 2D input. The combined inputs are then processed similarly with the sequence attention module, except for the additional 2D attention layer before the last convolution layer. The average of the outputs from the two modules is used as the final predicted contact map.

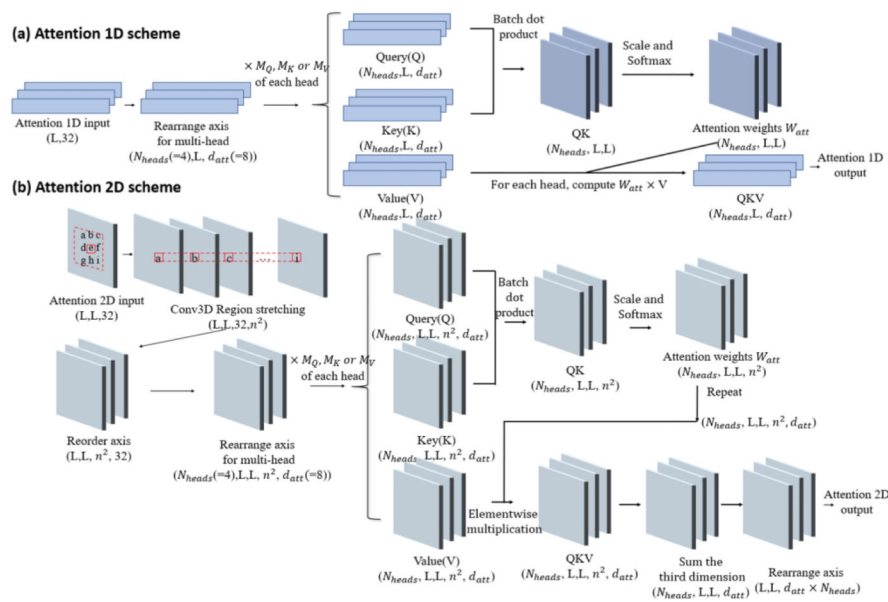


Figure 2. Schematic illustration of 1D and 2D attention mechanism. **(a)** The scheme for 1D attention mechanism. The input is first transformed into a vector of size (N_{heads}, L, d_{att}) for the efficient multi-headed attention implementation. For each head, the vector of size (L, d_{att}) is multiplied to three different trainable matrices of size (d_{att}, d_{att}) to generate Query(Q), Key(K), and Value(V). Q and K first go through a batch dot product operation, resulting in a new vector QK with size (N_{heads}, L, L) . QK is then scaled and normalized with Softmax function on the last axis, which becomes the attention score W_{att} . The product of $W_{att} \times V$ for each head becomes the 1D attention output. **(b)** The scheme for 2D attention mechanism. The 2D input is first transformed with a 3D convolution and becomes a stretched vector of size $(L, L, 32, n^2)$. It is then computed with the similar attention operation as the 1D attention scheme on the last axis.

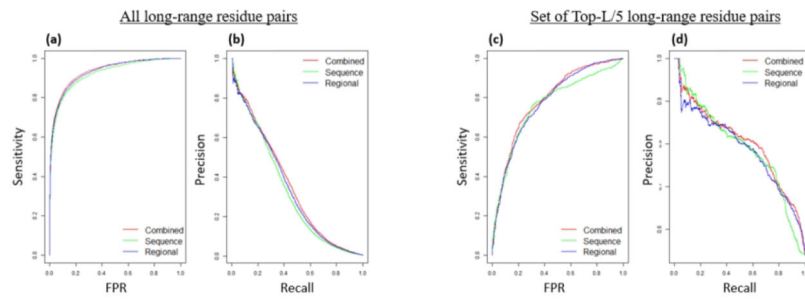


Figure 3. Prediction performance curves of the sequence attention model, regional attention model, and combined model. **(a)** ROC curve for all long-range contact predictions. **(b)** Precision-Recall curve for all long-range contact predictions. **(c)** ROC curve for all residue pairs that appear in the union of residue pairs from top-L/5 scores in any of the three models. **(d)** Precision-Recall curve for all residue pairs that appear in the top-L/5 scores in any of the three models.

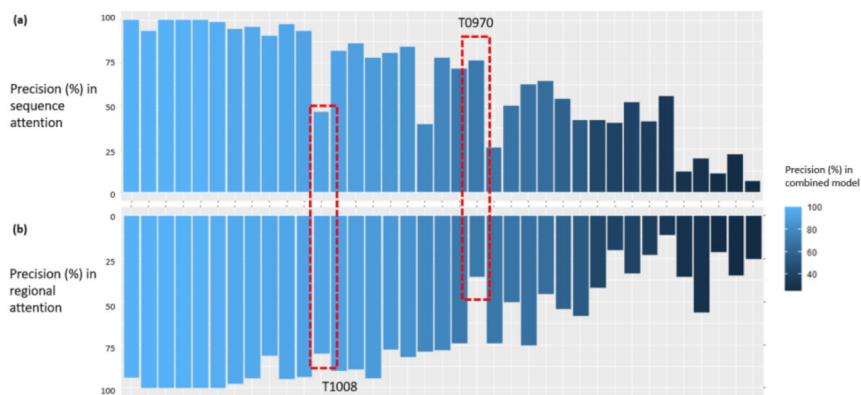


Figure 4. Comparison of the top-L/5 precision between sequence and regional attention module. The targets are arranged in the descending order of the top-L/5 precision in the combined model. **(a)** Precision scores from the sequence attention module. **(b)** Precision scores from the regional attention module. T1008 and T970 are two examples in which the two modules perform very differently.

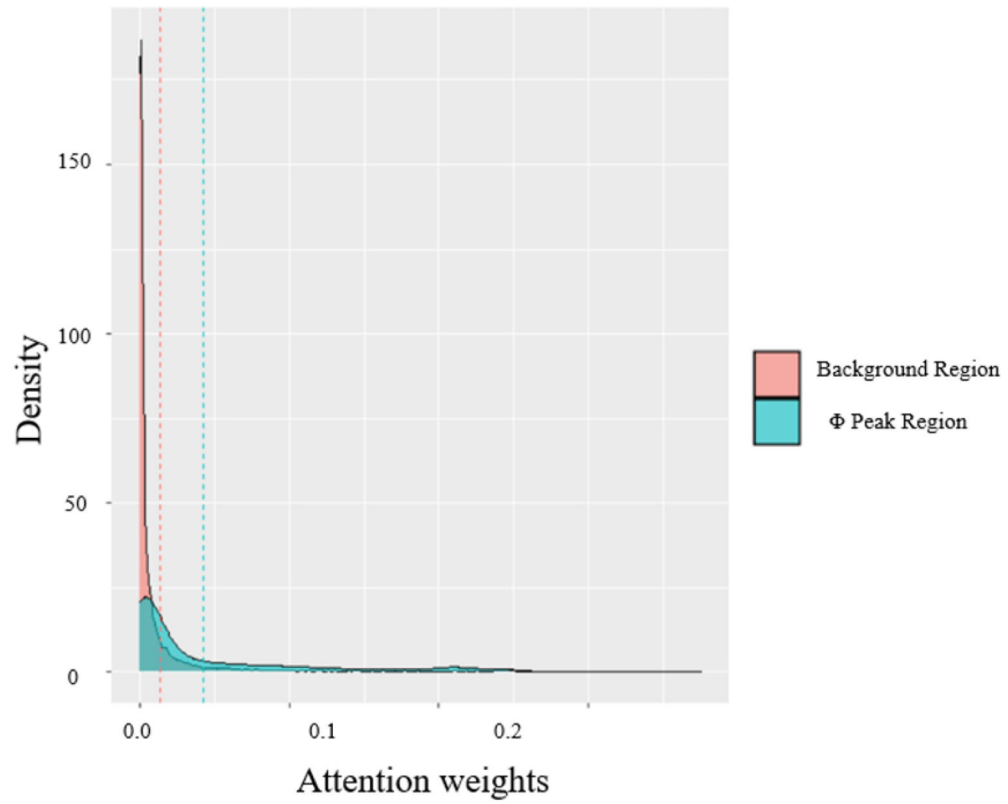


Figure 5. Comparison of attention scores from regions of the highest Φ -value peak and scores from the rest regions. The attention scores are averaged across all four heads.

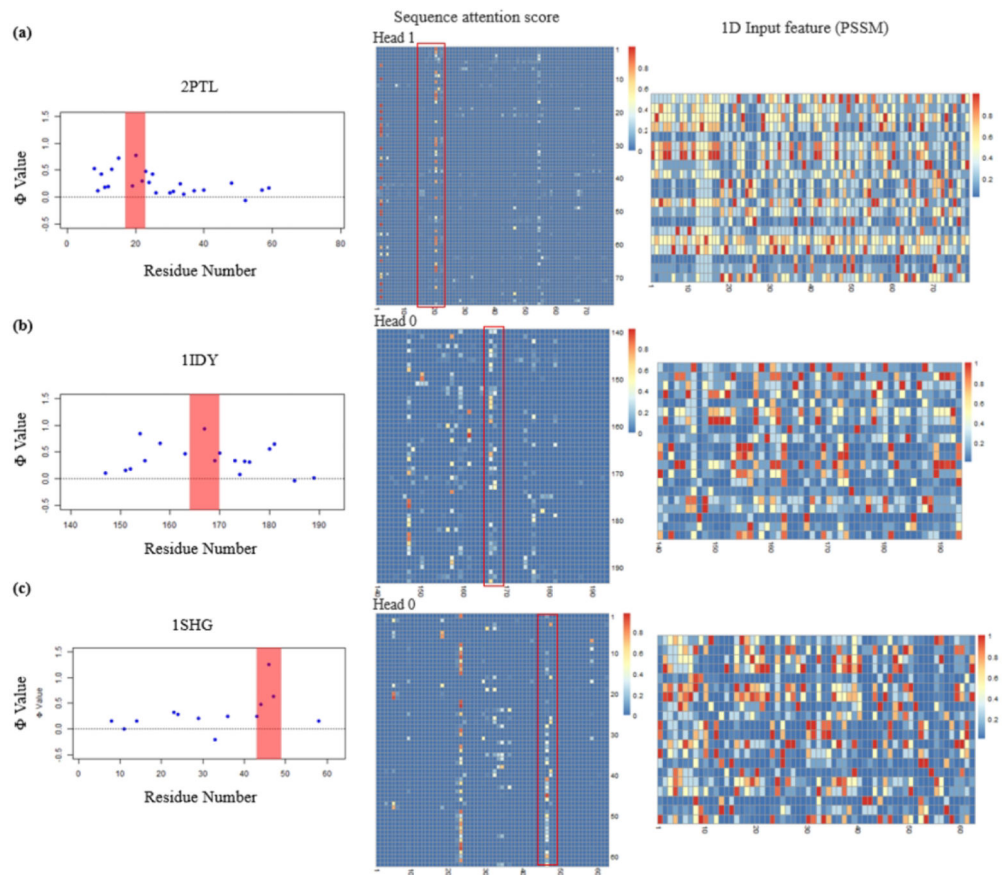


Figure 6. The visualization of Φ -values, sequence attention score and 1D input feature (PSSM) for three proteins 2PTL (a), 1IDY (b) and 1SHG (c)

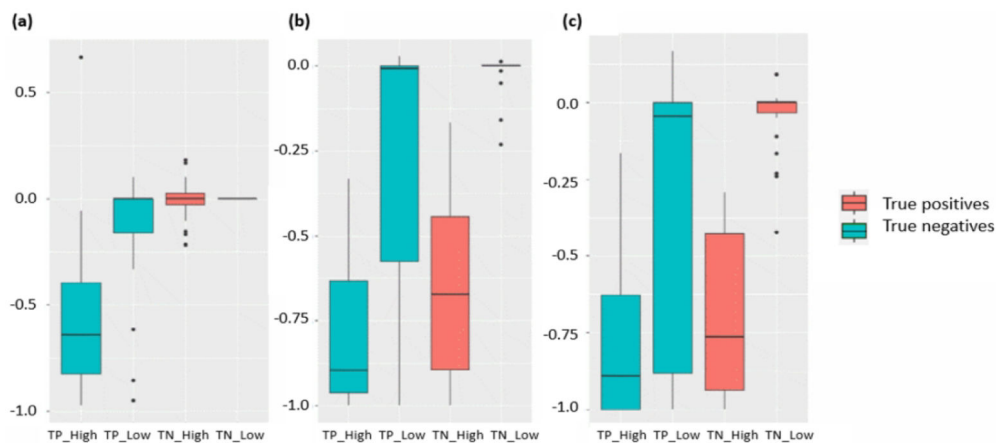


Figure 7.

Performance after permutation of different locations of the input. The Y-axis indicates the increase or decrease of top-L/5 precision scores after permutation. Here we choose locations that have the highest and lowest k attention scores as centers for permutated regions, where k is the number of true positives of each target. **(a)** Impact of permutated regions with size (1,1). **(b)** Impact of permutated regions with size (3,3). **(c)** Impact of permutated regions with size (5,5). TP_High: true positive predictions with high scores. TP_Low: true positive predictions with low scores. TN_High: true negative predictions with high scores. TP_Low: true negative predictions with low scores.

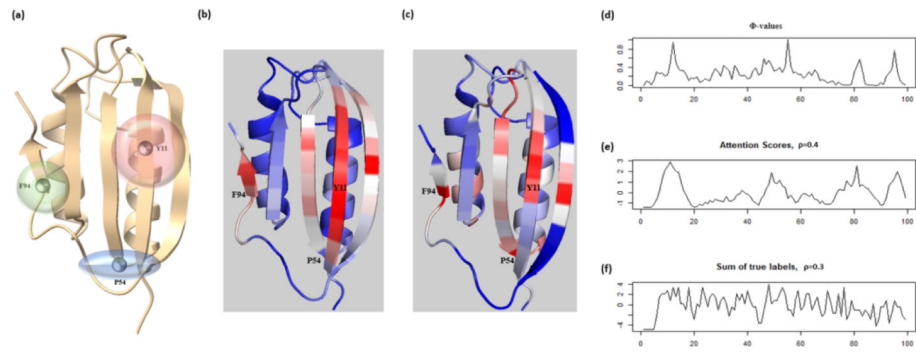


Figure 8.

Visualization and interpretation contact predictions of Human common-type acylphosphatase from the regional attention module. **(a)** The 3D model of AcP with the three highlighted critical residues in protein folding. The transparent spheres around the residues indicate their corresponding scopes in the contact networks. **(b)** The heatmap of regional attention scores shown on the 3D structure of AcP. **(c)** The heatmap of Φ -values shown on the 3D structure of AcP. **(d)**, **(e)**, **(f)** The Φ -values, attention scores and the count of true contacts for each residue plotted along the protein sequence. ρ : Pearson Correlation Coefficient.

Table 1.

Precision (%) of the top L/5, L/2 and L predicted long-range contacts on the CASP13 dataset

Model	Short-range			Medium-range			Long-range		
	Top-L/5	Top-L/2	Top-L	Top-L/5	Top-L/2	Top-L	Top-L/5	Top-L/2	Top-L
Sequence attention	58.26	41.51	27.95	63.10	45.87	32.51	64.46	52.13	39.82
Regional attention	61.08	41.95	28.38	65.46	48.00	33.65	67.32	54.15	40.96
Combined model	60.94	42.69	28.83	66.45	48.45	34.19	70.73	55.88	42.64
Baseline model	59.00	42.73	28.17	64.29	48.01	33.05	66.31	49.42	36.40

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2.

Comparison of the performance of the combined attention model with top 10 CASP13 methods.

long-range	Top-L/5	Top-L/2	Top-L
RaptorX-Contact	71.70	59.02	45.58
Combined Attention	70.73	55.88	42.64
TripletRes	65.97	55.34	42.65
ResTriplet	65.36	54.81	41.84
DMP	62.76	48.90	37.69
TripletRes_AT	60.77	52.02	40.13
RRMD	60.29	49.60	38.50
ZHOU-Contact	59.66	49.42	38.16
RRMD-plus	58.63	47.86	36.98
ResTriplet_AT	58.22	49.18	38.48
Zhang_Contact	58.07	49.58	39.21

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 3.

Comparison of the performance of different attention configurations. Bold numbers denotes the highest score achieved in each attention category.

ID	Model_type	Number_heads	Region	Top-L/5	Top-L/2	Top/L
1	Sequence	1	-	0.619499	0.460423	0.345457
2	Sequence	2	-	0.596136	0.453989	0.344475
3	Sequence	4	-	0.6356	0.474	0.3556
4	Regional	1	5	0.658928	0.479804	0.352311
5	Regional	2	5	0.664635	0.478009	0.354849
6	Regional	4	5	0.6708	0.4925	0.3675
7	Regional	4	3	0.660694	0.482968	0.35811

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript