Highlights

- Visual method achieves 67% sensitivity or recall (true positive rate) surpassing existing methods by over 35%.
- Significant improvement in detection rate by using dataset with naturally occurring class imbalance.
- Protein-peptide binding sites can be predicted using convolutional neural networks.
- Novel method for transforming protein data into image-like representations.

Graphical Abstract

Predicting protein-peptide binding sites with a Deep Convolutional Neural Network

Wafaa Wardah^{a,*}, Abdollah Dehzangi^b, Ghazaleh Taherzadeh^c, Mahmood A. Rashid^{e,d}, M. G. M. Khan^a, Tatsuhiko Tsunoda^{f,g,h,i}, Alok Sharma^{d,g,h,j,*}

^a*School of Computing, Information and Mathematical Sciences, Faculty of Science, Technology and Environment, The University of the South Pacific, Suva, Fiji*

^b*Department of Computer Science, Morgan State University, Baltimore, USA*

^c*Institute for Bioscience and Biotechnology Research, University of Maryland, USA*

^d*Institute for Integrated and Intelligent Systems, Griffith University, Queensland, Australia*

^e*Institute for Sustainable Industries and Liveable Cities, Victoria University Melbourne, Victoria, Australia*

^f*Department of Medical Science Mathematics, Medical Research Institute, Tokyo Medical and Dental University, Tokyo, Japan*

^g*Laboratory for Medical Science Mathematics, RIKEN Center for Integrative Medical Sciences, Yokohama, Japan*

^h*CREST, JST, Tokyo 113-8510, Japan*

ⁱ*Laboratory for Medical Science Mathematics, Department of Biological Sciences, Graduate School of Science, The University of Tokyo, Tokyo, Japan*

^j*School of Engineering and Physics, The University of the South Pacific, Suva, Fiji*

Abstract

Motivation: Interactions between proteins and peptides influence biological functions. Predicting such bio-molecular interactions can lead to faster disease prevention and help in drug discovery. Experimental methods for determining protein-peptide binding sites are costly and time-consuming. Therefore, computational methods have become prevalent. However, existing models show extremely low detection rates of actual peptide binding sites in proteins. To address this problem, we employed a two-stage technique - first, we extracted the relevant features from protein sequences and transformed them into images applying a novel method and then, we applied a convolutional neural network to identify the peptide binding sites in proteins.

Results: We found that our approach achieves 67% sensitivity or recall (true positive rate) surpassing existing methods by over 35%.

Keywords: protein-peptide binding, artificial intelligence, deep learning, convolutional neural network, protein sequence

*Corresponding authors

Email addresses: wafaa.wardah@usp.ac.fj (Wafaa Wardah), alok.fj@gmail.com (Alok Sharma)

1. Introduction

For the many roles that proteins play in and around cells, interacting with other molecules is known to be what enables most biological functionalities. To perform biological processes, proteins interact with a variety of molecular structures, such as nucleic acids (RNA and DNA) [47, 27], lipids, various small ligands [31] and other proteins. Some are covalent interactions including disulphide bonding and electron sharing, and others are weaker interactions including hydrogen bonds, hydrophobic interactions, Van der Waals forces and ionic interactions [45]. The presence of water molecules also plays a vital role in the interactions that occur [15]. While DNA repair, replication, gene expression and metabolism are known to be some of the vital cellular processes that protein interactions facilitate, studies have found that such interactions can also induce abnormal cellular behavior and disease such as cancers, where up to 40% of these interactions involve binding with relatively small peptides [24]. Therefore, analyzing protein-peptide interactions is necessary for understanding the molecular factors leading to various diseases [25, 20] and drug discovery [41]. Identifying the residues that are involved in these interactions and understanding the mechanisms that result in the binding of proteins and peptides are vital. In vivo methods currently used in this field include the Yeast two - hybrid screening method and affinity purification which involve high-throughput screening such as mass spectrometry and Nuclear Magnetic Resonance spectroscopy, which are expensive, labor-intensive and time-consuming. Through such experiments, protein data have been archived in repositories and are mostly accessible to the public. The first such repository was the Database of Interacting Protein (DIP) [46]. Other repositories that contain protein interaction related information include the Protein Data Bank (PDB) [2], BioLip [48] and Menthia [7].

Availability of protein data allowed machine learning techniques to be applied to the protein-peptide binding site prediction problem. Various methods were proposed to predict the peptide-binding sites of specific protein domains such as Major Histocompatibility Complex (MHC), PDZ (PSD-95, Discs-large, ZO-1), Src homology 2 (SH2) and Src homology 3 (SH3). In [9], researchers developed models MHC2SK and MHC2SKpan to computationally predict MHC binding peptides. They used Blosom62 for feature preparation [12] and proposed a model inspired by Spectrum RBF string kernel (SRBF) [38]. Another study [13] employed the Swiss-Prot database and used RBF kernel functions with Support Vector Machine (SVM) classifiers to predict peptide-binding sites of SH3 domains. In [19], researchers developed yet another SVM-based tool to predict protein-peptide binding sites of SH2 domains. Furthermore, PDZ-DockScheme was developed using a simulated annealing algorithm and rotamer optimization to predict protein-peptide bindings in PDZ domains [26]. A limitation of these methods is that they require known protein structures, while in reality, most protein structures are still unknown. Methods that are able to utilize protein sequence information to achieve reasonably accurate protein-peptide binding residues are a valuable contribution to the scientific community.

Figure 1: The diagram shows a simplified overview of the protein-peptide binding site prediction problem. The protein’s constituent amino acid sequence information is fed into the computational engine, which then produces the identified peptide binding sites.

The problem of protein-peptide interaction can be viewed as a binary classification problem, where in a protein chain, each residue can be classified into one of two classes: binding or non-binding (shown in Figure 1). Some techniques that have been used with this strategy include SVM, random forest and artificial neural networks (ANN). The models generally employ a sliding window to input the properties of the constituent residues along protein chains. The properties (or features) explored across literature include sequence, structural, evolutionary and physicochemical information [36, 37].

The common pipeline involves obtaining protein-ligand interaction data from BioLip [49], a semi curated database that derives protein data from the PDB [3]. Features are then selected, and the choice of computational technique is applied. In 2007, authors of SPPIDER [29] used multiple classifiers including SVM, ANN and linear discriminant analysis (LDA), where they studied the usefulness of the relative solvent accessibility (RSA) [42] feature. PSIVER [23] used inherent Naive Bayes classifier with kernel density estimation methods. LORIS [8] and SPRINGS [34] applied L1-regularized logistic regression and ANN methods to multiple input features. The method SPRINT is a server that makes sequence-based predictions of protein-peptide binding sites using SVM classifier [36]. The features used were sequence information in the form of 20 dimensional binary vectors, evolutionary information obtained from PSI-BLAST in the form of position specific scoring matrix (PSSM) [1], structural information obtained from the prediction tool SPIDER 2.0 [11] containing solvent accessible surface area (ASA) and secondary structure (SS), and physicochemical properties (steric parameter, hydrophobicity, volume, polarizability, isoelectric point, helix probability, and sheet probability) obtained from the AA Index database [17]. These features were fed through the SVM model, which produced a predicted label [39]. SPRINT-Str [37] was later developed by the same team to predict protein-peptide binding sites using the random forest classifier [6], employing additional protein 3-dimensional structure features and similar validation techniques. CRF-PPI [44] and SSWRF [43] are similar methods and also use SVM and random forest algorithms.

The problem that exists but gets ignored is that all current methods have extremely low rates of detecting actual binding sites. In most proteins, about

94% of the residues do not bind with peptides. If a tool predicted all residues as non-binding without running any computation at all, it would achieve an impressive 94% accuracy and 100% specificity score, but unfortunately would detect zero binding sites in the proteins. The utility of a protein-peptide binding site predictor lies in its ability to detect the actual binding sites, this measure is known as sensitivity, and the current methods are limited to very low sensitivity scores. The challenge is to develop a predictor whose classification is not skewed towards the non-binding sites and at the same time achieve a high binding site detection rate. In this paper, we have tried to address this problem with our proposed method, Visual, which uses a generally successful deep learning technique, convolutional neural networks (CNN), to predict the binding sites in proteins. Other methods have utilized CNNs for binding site identification, however, none have been used for the same objective. For example, DeepMHC [14] uses a 1-dimensional CNN model for predicting whether a given peptide will bind with a specific protein domain, MHC. They use 1-dimensional multi-channel one-hot vectors to represent the 13 amino acid long peptide sequence as input. On the contrary, Visual uses each individual protein amino acid as input to predict whether it will bind with a peptide. Another very successful tool, DeepSite [16] uses a deeper CNN model to identify pockets in proteins where ligands are likely to bind. It uses 3D images as input, which contributes to the high performance it achieves. However, this tool is for identifying binding sites for all ligands, and not specifically peptides. The overall high performance of CNNs in various domains as well as protein interaction is evident. The area to optimize now is the derivation of the protein information and representation so it can be used as CNN input. A recently developed technology, DeepInsight [33] is a general model that can be applied for non-image samples, like protein data. This is the first technique that applies three steps of element arrangement, feature extraction and classification. The element arrangement step is used to arrange an image suitable for CNNs. In DeepInsight, data is transformed into images by applying either KPCA (kernel principal analysis component) or t-SNE (t-distributed stochastic neighbor embedding) followed by convex hull algorithm. In this work, we have employed a simple yet novel method whereby extracted and calculated features of protein sequences, such as residue sequence and locality, sequence-based structure predictions, evolutionary information and physicochemical properties, are arranged into image-like representations that are then processed by a CNN. Visual algorithm detects over twice as many binding sites in the same dataset as previously published works.

2. Materials and methods

This section discusses the dataset used in this work, followed by a description of the methods used to transform the data and predict the protein-peptide binding sites in proteins.

2.1. Dataset

Protein-peptide binding data was extracted from BioLiP [49], where chains with less than 30 amino acid residues were said to be peptides. Redundant proteins that had more than 30% similarity were removed using the Blastclust toolkit [4]. From this set, 10% of the proteins were randomly selected and set aside as the independent test set, TS125. Another 10% of the proteins were randomly selected and used as a validation set, while the remaining proteins formed the training set. The resulting training, validation and independent test sets contained 1,004 proteins with 243,766 residues, 112 proteins with 22,823 residues and 125 proteins with 30,870 residues, respectively.

Each residue in the protein sequence is classified as either positive (binding) or negative (non-binding). All the subsets described above have a class ratio of approximately 17:1, this means, on average, there is about 1 peptide binding site in an 18 residue long protein segment. There are way more non-binding residues in a protein than those that are binding, so as a binary classification problem, this is a case of highly imbalanced class distribution.

2.2. Features

Features of the constituent amino acid residues were derived in various ways such that a detailed description of each site was available for the implemented model to classify. These features are discussed below:

- **Half sphere exposure (HSE)** is a measure of how buried an amino acid is in the protein 3D structure. The HSE values are calculated based on the contact numbers of the upward and downward hemispheres, as well as the pseudo $C\beta$ - $C\alpha$ bonds [10].
- A deep learning-based predicting tool, SPIDER2 [50], was used to obtain extended information about each residue. The tool has shown impressive prediction results, as noted in the literature.
 - **Secondary Structure (SS)** provides perspective into the local 3-dimensional conformation of the protein. Predicted values include probabilities for each of the three classes, α -helix, β -sheet and coil.
 - The **Accessible Surface Area (ASA)** describes the degree of solvent accessibility of a residue within a protein.
 - **Local backbone angles** include θ , τ , ϕ and ψ . These are torsion angles between contiguous residues and provide insight into the residue’s geometric relation to its locality.
- The **PSSM** is another feature that is available for proteins that has been widely used in the literature. The sequence-profiles were obtained from PSI-BLAST [1] using E-value threshold of 0.001 in three iterations to extract the 20-dimensional vector for each amino acid in the protein.

- **Physicochemical Properties** - The amino acid residues can also be represented with their physicochemical properties. These include steric parameter, hydrophobicity, volume, polarizability, isoelectric point, helix and sheet probability [17].

2.3. Method

CNN classifiers require images as inputs, however, the protein data used here is not in an image format. Therefore, we first transformed the data into image to feed in the network.

2.3.1. Input transformation

Various properties of the Amino Acid residues were used to create the image-like representations. The following sections describe how we transformed the protein data.

Feature vector

First, each residue's features, $F_1, \dots, F_6$, were combined to form a feature vector.

$$\begin{aligned}
F_1 &= \{HSE_1, HSE_2, HSE_3\} \\
F_2 &= \{SS_1, SS_2, SS_3\} \\
F_3 &= \{ASA\} \\
F_4 &= \{\theta, \tau, \phi, \psi\} \\
F_5 &= \{PSSM_{i,m}\} \text{ for all } m = 1, \dots, 20, \text{ where } m \text{ represents the } m^{th} \\
&\text{column of the } i^{th} \text{ amino acid} \\
F_6 &= \{PP_1, \dots, PP_7\}
\end{aligned}$$

So, if k^{th} protein sequence P_k has n residues; i.e., $P_k = \{R_1, R_2, \dots, R_n\}$, then $R_i = \{F_1, F_2, \dots, F_6\}$. Therefore, stacking the features horizontally results in a feature vector R_i of size 38, that contains information about the residue's structure, evolution and physicochemical properties.

Windowed segment

A sliding window approach is used to capture the locality of residues. Each residue R_i is represented by a segment S_{R_i} that is of a fixed length. We have used a window size of 7 as it is not too small as to not be effective and not too large as to hinder the computer hardware performance. The residue of concern R_i is in the center of adjacent residues, 3 upstream and 3 downstream, are also captured in S_{R_i} .

$$\text{Generally, } S_{R_i} = \{R_{i-3}, \dots, R_i, \dots, R_{i+3}\} \text{ for all } 4 \leq i < n - 3.$$

For residues that are at either edge of the protein sequence, $1 \leq i \leq 3$ and $n - 3 \leq i \leq n$, the missing side of the segment was augmented by mirroring the available residues. The expressions below describe how the missing residues (marked with *) were created. The start edge of the sequence is shown below in Figure 2:

Figure 2: (a) shows $S_{R_1} = \{R_{i+3}^*, R_{i+2}^*, R_{i+1}^*, R_i, R_{i+1}, R_{i+2}, R_{i+3}\}$ where $i = 1$, (b) shows $S_{R_2} = \{R_{i+2}^*, R_{i+1}^*, R_i, R_{i+1}, R_{i+2}, R_{i+3}, R_{i+4}\}$ where $i = 2$, and (c) shows $S_{R_3} = \{R_{i+1}^*, R_i, R_{i+1}, R_{i+2}, R_{i+3}, R_{i+4}, R_{i+5}\}$ where for $i = 3$.

The result was a set of 2-dimensional arrays that could be normalized and used in a CNN classifier as 7×38 pixel images. The features were normalized so that each value was in a range of 0 and 1. When viewed as a greyscale image, the lighter pixels represent high values, and the darker pixels represent low values (refer to Figure 3).

Figure 3: An example of an image-like input representing the center residue Serine (S), with window size = 7. Thus, a 7×38 image that can be classified with a CNN classifier. In order from left to right: 3 pixels represent the HSE values, 3 pixels represent the SS predicted probabilities, 1 pixel represents the ASA value, 4 pixels represent the backbone angles, 20 pixels represent the PSSM, and 7 pixels represent the Physicochemical properties of the amino acids.

2.3.2. Convolutional neural network - CNN

CNNs are deep neural networks that process data that come in the form of multiple arrays where the local values are so closely related that they form detectable motifs, like images. The concept of CNN arose from the workings of the biological visual system in humans. Given a field of view, the visual system scans patches of the field and learns to recognize the objects (feature maps) based on those observed patches. Similarly, here, we convert the protein features into visual data and train the CNN to learn the feature maps in the given synthesized images. CNN models involve complex matrix operations which demand high processing power. This work was achieved using the GTX1060Ti graphics card, programmed using PyTorch, a Python based deep learning platform [30]. The source code for the Visual model can be accessed online [40]. The Visual model consists of 2 sets of convolution layers, followed by a pooling layer and a fully connected layer. In the first convolution layer, 256 $[3 \times 3]$ kernels slide

over the $[7 \times 38]$ input image performing convolution operation resulting in 256 $[5 \times 36]$ convolved feature maps. These are further transformed with the application of the Rectified Linear Unit (ReLU) activation function ($y = \max(0, x)$), producing rectified feature maps. The second convolution layer contains 256 $[2 \times 2]$ kernels that perform convolution operation over the 256 $[5 \times 36]$ feature maps, resulting in 256 $[4 \times 35]$ convolved feature maps. The ReLU activation function is applied to these as well, followed by a pooling layer. In the pooling layer, the maximum values in every 2×2 patch of the feature maps are collected via a sliding window to form a more robust pooled feature map. These are then flattened into a 1×8704 feature vector, which is then fed through a fully connected layer, resulting in a 1×2 output that represents the two classes: non-binding and binding (refer to Figure 4).

Figure 4: The various transformations occurring throughout the CNN on a sample input are shown. The amino acid sequence is transformed into a $[7 \times 38]$ pixel image, which is fed into the CNN. The first convolution layer produces 256 5×36 matrices. The second convolution layer converts these into 256 $[4 \times 35]$ matrices. Next, the maxpooling layer transforms them into 256 $[2 \times 17]$ matrices, which are then flattened into a single $[1 \times 8704]$ vector. This is passed through a fully connected dense layer. This produces the final $[1 \times 2]$ vector, where the index 0 represents negative class and index 1 represents the positive class. In the sample shown, the predicted output is positive.

This model processed 128 samples per batch and an average loss per batch was calculated using the Cross-Entropy Loss function upon comparing the predicted outputs with the actual target labels. The internal weights of the network were adjusted using the Adam optimizer [18], which is an optimized variant of the gradient decent algorithm [32]. To avoid the problem of overfitting, early stopping technique was used to select the optimal number of epochs for training. It was found that the CNN model produced best validation results when trained for 18 epochs.

2.3.3. Bayesian Optimization of CNN hyperparameters

CNN models contain many hyperparameters that have varying effects on their overall performance. There are a few methods that aid in tuning these hyperparameters. Grid search is the strategy of trying out all possible values to arrive at the combination of hyperparameters that produces the best model. This is a useful method, however, it is very time-consuming. Another strategy is to try out values or options randomly to find the combination that produces the best model. Random search is also useful, but also relatively inefficient. A more apt strategy is to use some algorithm that produces values or choices to implement in the model. Bayesian optimization has shown promising results recently in finding hyperparameters for models most efficiently [35]. Rather than trying random values or trying every possible value, Bayesian optimization uses calculated values for configuring the model’s hyperparameters based on prior observations. The experiment was allowed to run for 350 iterations with different combinations of the two selected hyperparameters to achieve the best performing model. The first hyperparameter found using this method was the number of kernels in the two convolution layers of the CNN. The algorithm was given a list of options to select from $2^3, 2^4, 2^5, \dots, 2^{10}$. The second hyperparameter was the learning rate (alpha) used by the Adam optimizer when updating the weights of the network, where values ranged between 0.00001 and 0.001. The optimal hyperparameters were found at the 293rd iteration where *number of kernels* = $2^8 = 256$ and *Adam (alpha) learning rate* = 0.000091 (refer to Figure 5).

Figure 5: Flow diagram of the processes involved in (a). transforming the protein-binding data, (b). training CNN to achieve optimal internal state and (c). testing the final model using the test set TS125.

2.3.4. Performance evaluation

To effectively evaluate the performance of the method, the following values were calculated from the test outcome (confusion matrix):

- True positives (TP): the number of actual binding residues correctly predicted as binding sites.
- True negatives (TN): the number of actual non-binding residues correctly predicted as non-binding.
- False positives (FP): the number of actual non-binding residues incorrectly predicted as binding sites.

- False negatives (FN): the number of actual binding residues incorrectly predicted as non-binding sites.

Sensitivity is the measure of how well the actual binding sites are identified as binding sites. It is often called the recall, hit rate, or true positive rate (TPR). Sensitivity is a vital measure of performance in our case since being able to detect maximum residues in a protein that would bind with peptides can help in understanding protein interaction much better. Since binding sites are scarce in a protein, the rate at which the predictor method is able to detect them is vital. It is calculated by

$$Sensitivity = \frac{TP}{TP + FN} \quad (1)$$

Specificity is the ability of the predictor to correctly classify actual non-binding sites as such. In this case, most residues (about 94.4% of the test set TS125) do not bind with peptides. The specificity can be calculated by

$$Specificity = \frac{TN}{TN + FP} \quad (2)$$

Mathews correlation coefficient (MCC) is a score that is seen as a balanced measure that takes into account all 4 statistics from the confusion matrix. It can be calculated as

$$MCC = \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}} \quad (3)$$

An additional score, the area under the receiver operating characteristic (ROC) curve (AUC) is also obtained for the classifier. The ROC curve (see Figure 7) is a curve created by plotting the true positive rate (TPR or sensitivity) against the false positive rate (FPR or 1 - specificity).

3. Results and discussion

It is important to highlight the type of problem the protein-peptide binding site prediction case is. Although the problem can be dealt with as binary classification of the two classes (binding and non-binding sites for each residue in the protein sequence), it must be realized that the need to classify in the first place is to be able to correctly pick out the actual binding sites from the non-binding sites. Analyzing our model (Visual) with the unseen data set TS125, we have found that it is able to predict the class of each residue in a protein sequence with the highest sensitivity compared to any other tool. It is apparent that the protein-peptide binding sites predicted by Visual are quite close to the binding sites revealed by the experimental method (see Figure 6). Application areas, such as drug design, require reliable detection of binding sites. Visual is an attempt to improve this detection rate, and the results are positive with possibilities of further improvement.

Actual ——— KERLLDELTLGVARYMQSERCRRVICLVGAGISTSAGIPDFRSPSTGLYDNLEKYHLPYPEAIFEISYF }
 Predicted ——— KERLLDELTLGVARYMQSERCRRVICLVGAGISTSAGIPDFRSPSTGLYDNLEKYHLPYPEAIFEISYF }
 (55 - 125)
 { KKHPEFFALAKELYPGQFKPTICHYFMRLLKDKGLLLRCYTNIDTLERIAGLEQEDLVEAHGTFYTSH }
 { KKHPEFFALAKELYPGQFKPTICHYFMRLLKDKGLLLRCYTNIDTLERIAGLEQEDLVEAHGTFYTSH }
 (126 - 196)
 { CVSASCRHEYPLSWMKKIFSEVTPKCEDCQSLVKPDIVFQESLPARFFSCMQSDFLKVDLLVMGTSL }
 { CVSASCRHEYPLSWMKKIFSEVTPKCEDCQSLVKPDIVFQESLPARFFSCMQSDFLKVDLLVMGTSL }
 (197 - 267)
 { QVQPFASLISKAPLSTPRLLINKEKAGQSDPFLGMIMGLGGGMDFDSSKAYRDVAWLGECDQCLALAE }
 { QVQPFASLISKAPLSTPRLLINKEKAGQSDPFLGMIMGLGGGMDFDSSKAYRDVAWLGECDQCLALAE }
 (268 - 338)
 { LGWKKLEDLVREHASIDAQS }
 { LGWKKLEDLVREHASIDAQS }
 (339 - 361)

(a)

Figure 6: (a) shows the amino acid sequence of the protein 4l3oA. The sequence has been split into segments (55-125, 126-196, ...) to fit on the page. The upper row shows the actual peptide binding sites (magenta) while the lower row shows the peptide binding sites (magenta) as predicted by our method. (b) and (c) show computer-generated images of the protein 4l3oA, where (b) shows the actual binding sites in magenta and (c) shows the predicted binding sites in magenta.

Since about 94.4% of the residues do not bind with peptides and only 5.6% of them do, using the accuracy of the predictor as a basis for judging performance is misleading. The goal is to correctly classify the actual binding sites (which are only 5.6% of the data set) as such while keeping the number of falsely classified binding sites as low as possible. This trade-off between the true-positive and false-positive rates is depicted in the Receiver Operating Characteristics (ROC) curve shown in Figure 7. A method that predicts peptide binding sites randomly (poor predictor) will have a linear diagonal curve (dashed line) and the area under the ROC curve (AUC) will be 0.5, whereas the best predictor will be higher and have $AUC = 1$. Visual achieves $AUC = 0.73$, which is higher than all other methods except one, SPRINT-Str [37] (however, this method shows extremely low detection rate of peptide binding sites). Methods published previously have generally shown very low rates of detecting actual binding sites correctly. The highest sensitivity so far was by the method Peptimap which was able to correctly detect only 32% of the binding sites in the same test set

TS125. Our method is able to correctly detect 67% of binding sites in the test set TS125, the highest TPR so far achieved. Table 1 shows a comparison of the results of the top tools available for this problem.

Table 1: Comparison of different methods on the TS125 test set. The results stated were obtained from [37], where TS125 was used for comparison.

Methods		SEN	SPE	MCC	AUC
SPRINT-Str	[37]	0.24	0.98	0.29	0.78
SPRINT	[36]	0.21	0.96	0.20	0.68
Peptimap	[5]	0.32	0.95	0.27	0.63
Pepsite	[28]	0.18	0.97	0.20	0.61
PinUp	[22]	0.24	0.91	0.13	0.58
VisGrid	[21]	0.24	0.93	0.15	0.63
Visual	proposed method	0.67	0.68	0.17	0.73

Figure 7: ROC curve for this method, Visual, on the independent test set TS125. The curve portrays the performance of the method by plotting the True Positive Rate against the False Positive Rate.

4. Future work

There is room for improvement and especially in reducing the number of non-binding residues that get falsely classified as binding sites. The improvement may come about with better data preprocessing so that the image that

is fed into the CNN for classification is richer in specific information about the residue. First improvement may be achieved by arranging the protein features $F1, \dots, F6$ in an optimized order. More sophisticated feature engineering and deeper CNN topology, such as recently published DeepInsight [33], may produce better protein-peptide binding site prediction results. Secondly, improvement may be achieved by using more advanced computing environment such that window size greater than 7, and various other CNN topologies can be experimented with. Additionally, more work can be done in employing other types of deep-learning methods, such as RNN, to predict peptide binding sites in proteins.

5. Conclusion

A deep learning method, Visual, that can predict the peptide binding sites in a protein was proposed. The method is a 2 layer convolutional neural network that uses an image-like representation of the constituent amino acids to detect the binding sites in a protein. The binding site detection rate of Visual (67%) is over twice as high as previously published methods (32%). It can be concluded that protein data can be transformed into image-like data usable by CNN methods, and that CNN can be successfully optimized to achieve better results compared to the current methods.

Availability of codes

The code for Visual can be found on GitHub via this link <https://github.com/WafaaWardah/Visual> [40].

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

- [1] S. F. Altschul, T. L. Madden, A. A. Schffer, J. Zhang, Z. Zhang, W. Miller, and D. J. Lipman. Gapped blast and psi-blast: a new generation of protein database search programs. *Nucleic Acids Research*, 25:33893402, 1997.
- [2] H M Berman, J Westbrook, Z Feng, G Gilliland, T N Bhat, H Weissig, I N Shindyalov, and P E Bourne. The protein data bank. *Nucleic Acids Res*, 28:235–42, 2016.
- [3] Helen Berman, Kim Henrick, and Haruki Nakamura. Announcing the worldwide protein data bank. *Nature Structural and Molecular Biology*, 10:1545–9985, 2003.
- [4] Andreas Biegert, Christian Mayer, Michael Remmert, Johannes Soding, and Andrei N. Lupas. The mpi bioinformatics toolkit for protein sequence analysis. *Nucleic Acids Research*, 34, 2006.
- [5] Tanggis Bohnuud, George Jones, Ora Schueler-Furman, and Dima Kozakov. Detection of peptide-binding sites on protein surfaces using the pep-map server. *Methods in molecular biology*, 1561:11–20, 2017.
- [6] Leo Breiman. Random forests. *Machine Learning*, 45(1), Oct 2001.
- [7] Alberto Calderone, Luisa Castagnoli, and Gianni Cesareni. mentha: a resource for browsing integrated protein-interaction networks. *Nature Methods*, 10:690691, 2013.
- [8] Kaustubh Dhole, Gurdeep Singh, P. Pai, Priyadarshini, and Sukanta Mondal. Sequence-based prediction of proteinprotein interaction sites with l1-logreg classifier. *Journal of Theoretical Biology*, 348:47–54, 2014.
- [9] Linyuan Guo, Cheng Luo, and Shanfeng Zhu. Mhc2skpan: a novel kernel based approach for pan-specific mhc class ii peptide binding prediction. *BMC Genomics*, 14(5):S11, 2013.
- [10] Thomas Hamelryck. An amino acid has two sides: A new 2d measure provides a different view of solvent exposure. *Proteins: Structure, Function, and Bioinformatics*, 59(1):38–48, 2005.
- [11] Rhys Heffernan, Kuldip Paliwal, James Lyons, Abdollah Dehzangi, Alok Sharma, Jihua Wang, Abdul Sattar, Yuedong Yang, and Yaoqi Zhou. Improving prediction of secondary structure, local backbone angles, and solvent accessible surface area of proteins by iterative deep learning. *Scientific Reports*, 5(11476), 2015.
- [12] S Henikoff and J G Henikoff. Amino acid substitution matrices from protein blocks. *Proceedings of the National Academy of Sciences*, 89(22):10915–10919, 1992.

- [13] Tingjun Hou, Zheng Xu, Wei Zhang, William A. McLaughlin, David A. Case, Yang Xu, and Wei Wang. Characterization of domain-peptide interaction interface: A generic structure-based model to decipher the binding specificity of sh3 domains. *Mol Cell Proteomics.*, 8(4):639649, 2009.
- [14] Jianjun Hu and Zhonghao Liu. Deepmhc: Deep convolutional neural networks for high-performance peptide-mhc binding affinity prediction. *bioRxiv*, 2017.
- [15] Joel Janin. Wet and dry interfaces: the role of solvent in proteinprotein and proteindna recognition. *Structure*, 7:277–279, 1999.
- [16] J Jimenez, S Doerr, G Martinez-Rosell, A S Rose, and G De Fabritiis. DeepSite: protein-binding site predictor using 3D-convolutional neural networks. *Bioinformatics*, 33(19):3036–3042, 05 2017.
- [17] Shuichi Kawashima, Hiroyuki Ogata, and Minoru Kanehisa. Aaindex: Amino acid index database. *Nucleic Acids Research*, 27(1):368–369, 1999.
- [18] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014.
- [19] Kousik Kundu, Fabrizio Costa, Michael Huber, Michael Reth, and Rolf Backofen. Semi-supervised prediction of sh2-peptide interactions from imbalanced high-throughput data. *PLOS ONE*, 8(5):1–15, 05 2013.
- [20] U Kuzmanov and A Emili. Protein-protein interaction networks: probing disease mechanisms using model systems. *Genome Med.*, 5:37, 2013.
- [21] B Li, S. Turuvekere, M. Agrawal, D. La, K. Ramani, and D. Kihara. Characterization of local geometry of protein surfaces with the visibility criterion. *Proteins*, pages 670–83, 2008.
- [22] Shide Liang, Chi Zhang, Eduardo Liu, Song, and Yaoqi Zhou. Protein binding site prediction using an empirical scoring function. *Nucleic acids research*, pages 3698–707, 2006.
- [23] Y Murakami and K Mizuguchi. Applying the naive bayes classifier with kernel density estimation to the prediction of protein-protein interaction sites. *Bioinformatics*, 26(15):1841–8, 2010.
- [24] Victor Neduva, Rune Linding, Isabelle Su-Angrand, Alexander Stark, Federico de Masi, Toby J Gibson, Joe Lewis, Luis Serrano, and Robert B Russell. Systematic discovery of new recognition peptides mediating protein interaction networks. *PLOS Biology*, 3(12), 11 2005.
- [25] R K Nibbe, S A Chowdhury, M Koyuturk, R Ewing, and M R Chance. Protein-protein interaction networks and subnetworks in the biology of disease. *Wiley Interdiscip Rev Syst Biol Med.*, 3:357–67, 2011.

- [26] Masha Y. Niv and Harel Weinstein. A flexible docking procedure for the exploration of peptide binding selectivity to known structures and homology models of pdz domains. *Journal of the American Chemical Society*, 127(40):14072–14079, 2005. PMID: 16201829.
- [27] Zhenling Peng and Lukasz Kurgan. High-throughput prediction of rna, dna and protein binding regions mediated by intrinsic disorder, 2015.
- [28] Evangelia Petsalaki, Alexander Stark, Eduardo Garca-Urdiales, and Robert B. Russell. Accurate prediction of peptide binding sites on protein surfaces. *PLOS Computational Biology*, 5(3):1–10, 03 2009.
- [29] Aleksey Porollo and Jaroslaw Meller. Prediction-based fingerprints of protein-protein interactions. *Proteins*, 66(3):630–45, 2007.
- [30] PyTorch. An open source machine learning framework that accelerates the path from research prototyping to production deployment. <https://pytorch.org/> retrieved on 15/3/2020.
- [31] D B Roche, D A Brackenridge, and L J McGuffin. Proteins and their interacting partners: An introduction to protein-ligand binding site prediction methods. *Int J Mol Sci.*, 16(12):29829–42, 2015.
- [32] Sebastian Ruder. An overview of gradient descent optimization algorithms. *ArXiv*, abs/1609.04747, 2016.
- [33] Alok Sharma, Edwin Vans, Daichi Shigemizu, Keith A. Boroevich, and Tatsuhiko Tsunoda. Deepinsight: A methodology to transform a non-image data to an image for convolution neural network architecture. *Scientific Reports*, 9, 2019.
- [34] G Singh, K Dhole, PP Pai, and S Mondal. Springs: Prediction of protein-protein interaction sites using artificial neural networks. *J Proteomics Computational Biol*, 1(1):7, 2014.
- [35] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. Practical bayesian optimization of machine learning algorithms. *Proceedings of the 25th International Conference on Neural Information Processing Systems - Volume 2*, pages 2951–2959, 2012.
- [36] Ghazaleh Taherzadeh, Yuedong Yang, Tuo Zhang, Alan Wee-Chung Liew, and Yaoqi Zhou. Sequence-based prediction of proteinpeptide binding sites using support vector machine. *Journal of Computational Chemistry*, 2016.
- [37] Ghazaleh Taherzadeh, Yaoqi Zhou, Wee-Chung Liew, Alan, and Yuedong Yang. Structure-based prediction of protein-peptide binding regions using random forest. *Structural bioinformatics*, pages 1–8, 2017.
- [38] Nora C. Toussaint, Christian Widmer, Oliver Kohlbacher, and Gunnar Ratsch. Exploiting physico-chemical properties in string kernels. *BMC Bioinformatics*, 11(8):S7, Oct 2010.

- [39] Vladimir Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, 2000.
- [40] Github repository. The source code is accessible at <https://github.com/WafaaWardah/Visual>.
- [41] Patrick Vlieghe, Vincent Lisowski, Jean Martinez, and Michel Khrestchatisky. Synthetic therapeutic peptides: science and market. *Drug Discovery Today*, 15(1/2), October 2009.
- [42] M Wagner, R Adamczak, A Porollo, and J Meller. Linear regression models for solvent accessibility prediction in proteins. *J Comput Biol.*, 12:355–69, 2005.
- [43] Z Wei, K Han, J Yang, H Shen, and D Yu. Proteinprotein interaction sites prediction by ensembling svm and sample-weighted random forests. *Neurocomputing*, 193:201212, 2016.
- [44] Z Wei, J Yang, H Shen, and D Yu. A cascade random forests algorithm for predicting protein-protein interaction sites. *IEEE Transactions on NanoBioscience*, 14(7):746–760, 2015.
- [45] Jukka Westermarck, Johanna Ivaska, and Garry L Corthals. Identification of protein interactions involved in cellular signaling. *Molecular and cellular proteomics*, 12:1752–63, 2013.
- [46] Ioannis Xenarios and David Eisenberg. Protein interaction databases. *Current Opinion in Biotechnology*, 12(4):334–339, 2001.
- [47] J Yan, S Friedrich, and L Kurgan. A comprehensive comparative review of sequence-based predictors of dna- and rna-binding residues. *Brief Bioinform.*, 17:88–105, 2016.
- [48] J Yang, A Roy, and Y Zhang. Biolip: a semi-manually curated database for biologically relevant ligand-protein interactions. *Nucleic Acids Res*, 41:1096–103, 2012.
- [49] Jianyi Yang, Ambrish Roy, and Yang Zhang. Biolip: a semi-manually curated database for biologically relevant ligandprotein interactions. *Nucleic Acids Research*, 41:D1096D1103, 2013.
- [50] Yuedong Yang, Rhys Heffernan, Kuldip Paliwal, James Lyons, Abdollah Dehzangi, Alok Sharma, Jihua Wang, Abdul Sattar, and Yaoqi Zhou. Spider2: A package to predict secondary structure, accessible surface area, and main-chain torsional angles by deep neural networks. *Prediction of Protein Secondary Structure*, pages 55–63, 2014.

Author Credit *JTB-D-19-01085*

Wafaa Wardah: Conceptualization, Methodology, Software, Writing - Original Draft

Abdollah Dehzangi: Data Curation

Ghazaleh Taherzadeh: Writing - Original Draft

Mahmood Rashid: Writing - Original Draft

M. G. M. Khan: Formal analysis

Tatsuhiko Tsunoda: Formal analysis

Alok Sharma: Conceptualization, Methodology, Validation, Writing - Review & Editing,
Supervision