

Protein Secondary Structure Prediction Using Cascaded Convolutional and Recurrent Neural Networks

Zhen Li, Yizhou Yu

Department of Computer Science, The University of Hong Kong
zli@cs.hku.hk, yizhouy@acm.org

Abstract

Protein secondary structure prediction is an important problem in bioinformatics. Inspired by the recent successes of deep neural networks, in this paper, we propose an end-to-end deep network that predicts protein secondary structures from integrated local and global contextual features. Our deep architecture leverages convolutional neural networks with different kernel sizes to extract multiscale local contextual features. In addition, considering long-range dependencies existing in amino acid sequences, we set up a bidirectional neural network consisting of gated recurrent unit to capture global contextual features. Furthermore, multi-task learning is utilized to predict secondary structure labels and amino-acid solvent accessibility simultaneously. Our proposed deep network demonstrates its effectiveness by achieving state-of-the-art performance, i.e., **69.7%** Q8 accuracy on the public benchmark CB513, **76.9%** Q8 accuracy on CASP10 and **73.1%** Q8 accuracy on CASP11. Our model and results are publicly available¹.

1 Introduction

Accurately and reliably predicting structures, especially 3D structures, from protein sequences is one of the most challenging tasks in computational biology, and has been of great interest in bioinformatics [Ashraf and Yaohang, 2014]. Structural understanding is not only critical for protein analysis, but also meaningful for practical applications including drug design [Noble *et al.*, 2004]. Understanding protein secondary structure is a vital intermediate step for protein structure prediction as the secondary structure of a protein reflects the types of local structures (such as 3_{10} -helix and β -bridge) present in the protein. Thus, an accurate secondary structure prediction significantly reduces the degree of freedom in the tertiary structure, and can give rise to a more precise and high resolution protein structure prediction [Ashraf and Yaohang, 2014; Zhou and Troyanskaya, 2014; Wang *et al.*, 2011].

The study of protein secondary structure prediction dates back to 1970s. In the 1970s, statistical models

were frequently used to analyze the probability of specific amino acids appearing in different secondary structure elements [Chou and Fasman, 1974]. The Q3 accuracy, i.e., the accuracy of three-category classification: helix (H), strand (E) and coil (C), of these models was lower than 60% due to inadequate features. In the 1990s, significant improvements were achieved by exploiting the evolutionary information of proteins from the same structural family [Rost and Sander, 1993] and position-specific scoring matrices [Jones, 1999]. During this period, the Q3 accuracy exceeded 70% by taking advantage of these features. However, progress stalled when it came to the more challenging 8-category classification problem, which needs to distinguish among the following 8 categories of secondary structure elements: 3_{10} -helix (G), α -helix (H), π -helix (I), β -strand (E), β -bridge (B), β -turn (T), bend (S) and loop or irregular (L) [Zhou and Troyanskaya, 2014; Yaseen and Li, 2014]. In the 21-st century, various machine learning methods, especially artificial neural networks, have been utilized to improve the performance, e.g., SVMs [Sujun and Zhirong, 2001], recurrent neural networks (RNNs) [Pollastri *et al.*, 2002], probabilistic graphical models such as conditional neural fields combining CRFs with neural networks [Wang *et al.*, 2011], generative stochastic networks [Zhou and Troyanskaya, 2014].

It is well known that local contexts are critical for protein secondary structure prediction. Specifically, the secondary structure category information of the neighbours of an amino acid are the most effective features for classifying the secondary structure this amino acid belongs to. For instance, in Fig.1, the 18th to 21th amino acids in PDB 154L [Simpson and Morgan, 1983] (obtained from the publicly available protein data bank ²) are likely to be assigned the same secondary structure label given their neighbours' information. Convolutional neural networks (CNN) [LeCun *et al.*, 1998], a specific type of deep neural networks using translation-invariant convolutional kernels, can be applied to extracting local contextual features and have proven to be effective for many natural language processing (NLP) tasks [Yih *et al.*, 2011; Zhang *et al.*, 2015]. Inspired by their success in text classification, in this paper, CNNs with various kernel sizes are used to extract multiscale local contexts from a protein sequence.

On the other hand, long-range interdependency among dif-

¹<https://github.com/icemansina/IJCAI2016>

²<http://www.rcsb.org/pdb/explore/explore.do?structureId=154L>




Figure 1: The amino acid sequence and its corresponding 3-state secondary structure of PDB 154L with UniProtKB accession number (P00718), which consists of 185 residues.

ferent types of amino acids also holds vital evidences for the category of a secondary structure, e.g., a β -strand is steadi- ed by hydrogen bonds formed with other β -strands at a distance [Zhou and Troyanskaya, 2014]. For instance, also in Fig.1, the 4th and 60th amino acids can be determined to share the same secondary structure label given the disulphide bond annotated as link 1. Similar to CNNs, recurrent neural networks (RNNs) are another specific type of neural networks with loop connections. They are designed to capture dependencies across a distance larger than the extent of local contexts. In previous work [Sepideh *et al.*, 2010], RNN models could not perform well on protein secondary structure prediction partially due to the difficulty to train such models. Fortunately, RNNs with gate and memory structures, including long short term memory (LSTM) [Hochreiter and Schmidhuber, 1997], gate recurrent units (GRUs) [Cho *et al.*, 2014a], and JZ3 structure [Jozefowicz *et al.*, 2015], can artificially learn to remember and forget information by using specific gates to control the information flow. In this paper, we exploit bidirectional gate recurrent units (BGRUs) to capture long-range dependencies among amino acids from the same protein sequence.

In summary, the main contributions of this paper are as follows.

- We propose a novel deep convolutional and recurrent neural network (DCRNN) for protein secondary structure prediction. This deep network consists of a feature embedding layer, multiscale CNN layers for local context extraction, stacked bidirectional RNN layers for global context extraction, fully connected and softmax layers for final joint secondary structure and solvent accessibility classification. Experimental results on the CB6133 dataset, the public CB513 benchmark, and the recent CASP10 and CASP11 datasets demonstrate that our proposed deep network outperforms existing methods and achieves state-of-the-art performance.
- To our knowledge, it is the first time to apply bidirectional GRU layers to secondary protein structure predic-

tion. An ablation study indicates they form the most important component of our deep neural network.

The rest of this paper is organized as follows. In Section 2, we introduce our proposed end-to-end deep model in detail. We present implementation details, experimental results and an ablation study in Section 3. Section 4 concludes this paper with final remarks.

2 Network Architecture

As illustrated in Fig. 2, our deep convolutional and recurrent neural network (DCRNN) for protein secondary structure prediction consists of four parts, one feature embedding layer, multiscale convolutional neural network (CNN) layers, three stacked bidirectional gated recurrent unit (BGRU) layers and two fully connected hidden layers. The input to our deep network carries two types of features of a protein amino acid sequence, sequence features and profile features. The feature embedding layer is responsible for transforming sparse sequence feature vectors into denser feature vectors in a new feature space. The embedded sequence features and the original profile features are fed into multiscale CNN layers with different kernel sizes to extract multiscale local contextual features. The concatenated multiscale local contexts flow into three stacked BGRU layers, which capture global contexts. On top of the cascaded CNN and BGRU layers, there are two fully connected hidden layers taking concatenated local and global contexts as input. The output from the second fully connected layer with softmax activation is fed into the output layer, which performs 8-category secondary structure and 4-category solvent accessibility classification.

2.1 Feature Embedding

For better understanding, protein secondary structure prediction can be formulated as follows. Given an amino acid sequence $X = x_1, x_2, \dots, x_T$, we need to predict the secondary structure label of every amino acid, $S = s_1, s_2, \dots, s_T$, where $x_i (\in \mathbb{R}^n)$ is an n -dimensional feature vector corresponding to the i -th amino acid, and s_i is an 8-state secondary structure label. In this paper, the input feature sequence X is decomposed into two parts, one is a sequence of 21-dimensional feature vectors encoding the types of the amino acids in the protein and the other is a sequence of 21-dimensional profile features obtained from the PSI-BLAST [Altschul *et al.*, 1997] log file and rescaled by a logistic function [Jones, 1999]. Note that each feature vector in the first sequence is a sparse one-hot vector, i.e., only one of its 21 elements is non-zero, while a profile feature vector has a dense representation. In order to avoid the inconsistency of feature representations, we adopt an embedding operation from natural language processing to transform sparse sequence features to a denser representation [Mesnil *et al.*, 2015]. This embedding operation is implemented as a feedforward neural network layer with an embedding matrix, $W_{emb} \in \mathbb{R}^{21 \times D_{emb}}$, that maps a sparse 21-dimensional vector into a denser D_{emb} -dimensional vector. In this paper, we empirically set $D_{emb} = 50$, and initialize the embedding matrix with random numbers. The embedded sequence feature

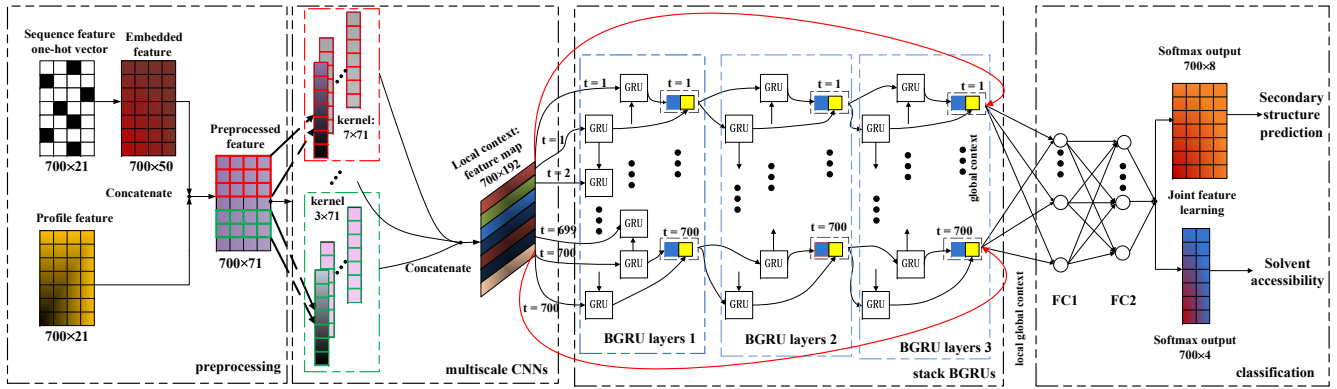


Figure 2: Our end-to-end deep convolutional and recurrent neural network (DCRNN) for predicting protein secondary structures. The input consists of sequence features and profile features. Through feature embedding and concatenation, preprocessed features are fed into the multiscale CNN layer, where multiple kernel sizes are used to extract multiscale local features. The concatenated multiscale features as local contexts flow into three stacked BGRU layers for capturing global contexts. On top of the stacked BGRU layers, two fully connected hidden layers are used for multi-task joint classification.

vector is concatenated with the profile feature vector before being fed into multiscale CNN layers.

2.2 Multiscale CNNs

As illustrated in Fig. 2, the second component of our deep network is a set of multiscale convolutional neural network layers. Given the amino acid sequence with embedded and concatenated features

$$\tilde{X} = [\tilde{x}_1, \tilde{x}_2, \dots, \tilde{x}_T], \quad (1)$$

where $\tilde{x}_i \in \mathbb{R}^m$ is the preprocessed feature vector of the i -th amino acid. To model local dependencies of adjacent amino acids, we leverage CNNs with a sliding window and Rectified Linear Unit (ReLU) [Nair and Hinton, 2010] to extract local contexts.

$$\tilde{l}_i = F * \tilde{x}_{i:i+f-1} = \text{ReLU}(w \cdot \tilde{x}_{i:i+f-1} + b),$$

where $F \in \mathbb{R}^{f \times m}$ is a convolutional kernel, f is the extent of the kernel along the protein sequence and m is the feature dimensionality at individual amino acids ($m = 71$ in this paper), b is the bias term and ‘ReLU’ is the activation function. The kernel goes through the full input sequence and generates a corresponding output sequence, $\tilde{L} = [\tilde{l}_1, \tilde{l}_2, \dots, \tilde{l}_T]$,

where each $\tilde{l}_i \in \mathbb{R}^q$ has q channels ($q = 64$ in this paper). Since an amino acid is sometimes affected by other residues at a relative large distance, e.g., two residues have interaction at a distance of 11 in the disulphide bond labeled as link 2 in Fig. 1, multiscale CNN layers with different kernel sizes are used to obtain multiple local contextual feature maps. In this paper, we use three CNN layers with $f = 3, 7,$ and 11 . This results in three feature maps $\tilde{L}_1, \tilde{L}_2, \tilde{L}_3$. These multiscale features are concatenated together as local contexts $L = \text{concatenate}\{\tilde{L}_1, \tilde{L}_2, \tilde{L}_3\}$.

2.3 BGRUs

In addition to local dependencies, long-range dependencies, such as line 1 in Fig. 1, also widely exist in amino acid sequences. Multiscale CNNs can only capture dependencies

among amino acids separated by a distance not larger than its maximum kernel size. To capture dependencies across a larger distance, we exploit bidirectional gate recurrent units (BGRUs).

Recurrent neural networks (RNNs) have a powerful capacity to deal with context-dependent sequences. However, training RNNs used to be difficult due to vanishing gradients [Bengio *et al.*, 1994]. Only in recent years RNNs with gated units, e.g., long short term memory (LSTM) and gate recurrent units (GRUs), became practically useful. In our network, GRUs [Cho *et al.*, 2014b] are used to capture global contexts because they achieve comparable performance with less parameters in comparison to LSTM [Jozefowicz *et al.*, 2015]. In the same notation as previous equations, the mechanism of a GRU (illustrated in Fig. 3) can be presented as follows if the input is (l_t, h_{t-1}) .

$$r_t = \text{sigm}(W_{lr} \cdot l_t + W_{hr} \cdot h_{t-1} + b_r), \quad (2)$$

$$u_t = \text{sigm}(W_{lu} \cdot l_t + W_{hu} \cdot h_{t-1} + b_u), \quad (3)$$

$$\tilde{h}_t = \tanh(W_{l\tilde{h}} \cdot l_t + W_{h\tilde{h}} \cdot (r_t \odot h_{t-1} + b_{\tilde{h}}), \quad (4)$$

$$h_t = u_t \odot h_{t-1} + (1 - u_t) \odot \tilde{h}_t, \quad (5)$$

where $r_t, u_t, \tilde{h}_t, h_t \in \mathbb{R}^k$ are respectively activations of the reset gate, update gate, internal memory cell, and GRU output if the number of hidden units is k ; $W_{lr}, W_{hr}, W_{lu}, W_{hu}, W_{l\tilde{h}}, W_{h\tilde{h}} \in \mathbb{R}^{3q \times k}$ are weight matrices; and $b_r, b_u, b_{\tilde{h}} \in \mathbb{R}^k$ are bias terms. In addition, \odot , sigm and tanh stand for element-wise multiplication, sigmoid and hyperbolic functions, respectively. Compared with LSTM, which has three gates (i.e., input gate, forget gate, output gate), one external memory cell state and one output state, a GRU only has two gates (update gate, reset gate) and one output state. It does not have the least important gate (output gate) in LSTM, and merges the input gate and the forget gate together to form the update gate and the reset gate, which control when information should be artificially remembered or forgotten. The total number of parameters in a GRU is only $3/4$ of that in LSTM [Jozefowicz *et al.*, 2015].




Figure 3: The internal structure of a gate recurrent unit (GRU).

The secondary structure label of an amino acid not only depends on the label of its preceding amino acid in the sequence, but also the label of its following amino acid. Thus, we use bidirectional GRUs, each of which consists of a forward GRU (from $t = 0$ to $t = T$) and a backward GRU (from $t = T$ to $t = 0$). The output from both forward and backward GRUs at time t are concatenated together to form the output from the bidirectional GRU (BGRU) at the same time. Furthermore, to enhance global information flow in our network, three BGRUs are stacked together with dropout to improve the performance. Note that the forward hidden state of the stacked BGRUs is calculated as $h_{forward,t}^z = \text{GRU}(h_t^{z-1}, h_{forward,t-1}^z)$, where z and h_t^{z-1} stand for the layer index and the concatenated output of the preceding layer in the stacked BGRUs. At last, the obtained local and global contexts are concatenated together as the input to the following layers.

2.4 Multi-Task Joint Feature Learning

Taking the interaction between different protein structure properties into consideration, we train our proposed model to generate two different but correlated types of results by performing joint feature learning in two shared fully connected hidden layers, as shown in the classification part of Fig. 2. Specifically, the output from our proposed model consists of the predicted sequences of secondary structure labels s_i and solvent accessibility labels a_i (four-category classification problem, i.e., absolute and relative solvent accessibility). Absolute and relative accessibility are judged by specific thresholds of the original and normalized solvent accessibility values computed by the DSSP program [Kabsch and Sander, 1983], and solvent accessibility is closely related to secondary structure prediction. According to the multi-task training method from [Ren *et al.*, 2015], l_2 -norms are added to the loss function as the regularization term. In addition, dropout [Srivastava *et al.*, 2014] is employed in the stacked BGRU layers and the penultimate layer to avoid overfitting. The joint loss function can be formulated as follows.

$$L(\{s_i\}, \{a_i\}) = \frac{1}{N} \sum_i L_s(s_i, s_i^*) + \lambda_1 \frac{1}{N} \sum_i L_a(a_i, a_i^*) + \lambda_2 \|\theta\|_2, \quad (6)$$

where $L_s(s_i, s_i^*) = -s_i^* \log(s_i)$ and $L_a(a_i, a_i^*) = -a_i^* \log(a_i)$ are respective loss functions for secondary struc-

ture prediction and solvent accessibility prediction, s_i and a_i are predicted probabilities of secondary structure labels and solvent accessibility labels respectively, s_i^* and a_i^* are ground-truth labels of secondary structure and solvent accessibility respectively, θ is the weight vector, and N is the number of residues.

3 Experimental Results

3.1 Datasets and Features

We use four publicly available datasets, CB6133 produced with PISCES CullPDB [Wang and Dunbrack, 2003], CB513 [Cuff and Barton, 1999]³, CASP10 [Kryshtafovych *et al.*, 2014] and CASP11 [Moult *et al.*, 2014], to evaluate the performance of our proposed deep neural network. CB6133 is a large non-homologous protein sequence and structure dataset, that has 6128 proteins, which include 5600 proteins (index 0 to 5599) for training, 256 proteins (index 5877 to 6132) for validation and 272 proteins (index 5605 to 5876) for testing. Note that the testing set of CB6133 is different from that in [Wang *et al.*, 2016]. CB513 is a public benchmark dataset used for testing only. Since there exists redundancy between CB513 and CB6133, a smaller filtered version of CB6133 is formed by removing sequences in CB6133 that have over 25% similarity with some sequence in CB513, as done in [Zhou and Troyanskaya, 2014; Wang *et al.*, 2016]. The filtered CB6133 dataset has 5534 proteins, which can all be used for training if CB513 is used as the testing set. CASP10 and CASP11 contain 123 and 105 domain sequences, respectively. They are used for testing the performance of our proposed network trained on the filtered CB6133 dataset. The performance of secondary structure prediction is measured by Q8 accuracy.

Every protein sequence in the aforementioned datasets has 55 channels of data per residue. Among the 55 channels, 21 channels are used for the sequence feature, which specifies the category of the amino acid, 21 channels for the sequence profile (PSSM rescaled through a logistic function, PSSM calculated by PSI-BLAST against the UniRef90 database with E-value threshold 0.001 and 3 iterations), 8 channels for secondary structure category labels, 2 channels for solvent accessibility labels (obtained by the DSSP program through 3D PDB). Although additional features could be considered to further improve performance, we focus on network architecture in this paper. Note that there exist 3 mask channels after the sequence feature, the sequence profile and the secondary structure label. For the convenience of subsequent processing and implementation, the length of all protein amino acid sequences in these datasets are normalized to 700. Sequences longer than 700 are truncated and those shorter than 700 are padded with zeros. Most sequences are shorter than 700.

3.2 Implementation Details

In our experiments, multiscale CNN layers with kernel size 3, 7, and 11 are used to extract local contexts given the window size of “a biological word” in [Asgari and Mofrad, 2015] as a reference. The obtained 3 feature maps, each having 64

³<http://www.princeton.edu/~jzthree/datasets/ICML2014/>

channels, are concatenated together as the local contextual feature vector. Each of the three stacked BGRU layers has 600 hidden units. They take the local contexts as the input. The output from the BGRU layers is regularized with dropout ($= 0.5$) to avoid overfitting. The local and global contexts obtained from multiscale CNN layers and BGRU layers are concatenated together and fed to two fully connected layers with ReLU activation. We set $\lambda_1 = 1, \lambda_2 = 0.001$ for balancing the two joint learning tasks and the regularization term.

We also exploit bagging to obtain ensemble models. According to the standard bagging algorithm, for each weak model, we randomly choose 512 (about 10%) proteins from the original training set to form the validation set and the remaining training samples form the training set. Our ensemble model consists of 10 independently trained weak models. Early stopping is used during training. Specifically, when the F1 score on the validation set is not increasing for 10 epochs, we decrease the learning rate by a factor of 2. Once the learning rate is smaller than a predefined threshold, we stop training, test the model obtained after every epoch on the validation set, and choose the one with the best performance on the validation set as our trained model.

Our code is implemented in Theano [Bastien *et al.*, 2012; Bergstra *et al.*, 2010], a publicly available deep learning software⁴, on the basis of the Keras [Chollet, 2015] library⁵. Weights in our neural networks are initialized using the default setting in Keras. We train all the layers in our deep network simultaneously using the Adam optimizer [Kingma and Ba, 2014]. The batch size is set to 128. The entire deep network is trained on a single NVIDIA GeForce GTX TITAN X GPU with 12GB memory. It takes about one day to train our deep network without early stopping while it only takes 6 hours if we take advantage of early stopping. In the testing stage, one protein takes 5ms on average.

3.3 Performance

We evaluate the overall performance of our deep network (DCRNN) by performing three sets of experiments. In the first set of experiments, we perform both training and testing on the original CB6133 dataset. In the second set, we perform training on the filtered CB6133 dataset and testing on the CB513 benchmark. In the third set, we still perform training on the filtered CB6133 dataset but performance is measured on the recent CASP10 and CASP11 datasets.

Training and Testing on CB6133

Our model trained using the training set of CB6133 achieves $73.2 \pm 0.6\%$ Q8 accuracy, which defines a new state of the art and is 1.1% higher than the previous best result obtained with GSN [Zhou and Troyanskaya, 2014], and 76.1% solvent accessibility accuracy on the testing set of CB6133. We did not compare against the result from [Wang *et al.*, 2016] because the testing set used in [Wang *et al.*, 2016] is different from the testing set defined in CB6133, and is not publicly available either. In Table 1, we compare our overall Q8 performance as well as the performance on individual secondary structure labels with the previous best results achieved by GSN [Zhou

Table 1: Classification precision and recall of individual secondary structure labels on the testing set of CB6133. Our results (DCRNN) are compared with GSN. Frequencies of secondary structure labels in the training set are given in the last column. Boldface numbers indicate best performance.

	Precision		Recall		Frequency
	DCRNN	GSN	DCRNN	GSN	
Q8	0.732	0.721			
H	0.878	0.828	0.927	0.935	0.344
E	0.792	0.748	0.862	0.823	0.217
L	0.589	0.541	0.662	0.633	0.193
T	0.577	0.548	0.572	0.506	0.113
S	0.518	0.423	0.275	0.159	0.083
G	0.434	0.496	0.311	0.133	0.039
B	0.596	0.500	0.049	0.001	0.010
I	0.000	0.000	0.000	0.000	0.000

and Troyanskaya, 2014]. Apparently, our proposed model achieves higher precision and recall on almost all individual labels. We believe that the better performance is not only due to the power of the neural networks we use, but also owing to the power of the integrated local and global contexts. Specifically, on the four labels (H, E, L and T) with high frequency, our model achieves better performance due to the fact that our model with millions of parameters has higher representation capacity. Nevertheless, our model also performs better than previous models on low frequency labels most likely because of the integrated local and global contexts, which form the core contribution of this paper.

Training on Filtered CB6133 and Testing on CB513

We have also performed validation on the public CB513 benchmark using a model trained on the filtered CB6113 dataset, whose 5534 proteins do not include any sequences with more than 25% similarity with proteins in CB513. Our single trained model achieves $69.4 \pm 0.5\%$ Q8 accuracy, which is 1.1% higher than the previous state of the art achieved by DeepCNF [Wang *et al.*, 2016] on the same training and testing sets, and 76.8% solvent accessibility accuracy on the validation set as ground-truth solvent accessibility labels are unavailable on CB513. We also compare our model against other existing methods (e.g., CNF [Wang *et al.*, 2011], SSpro8 [Pollastri *et al.*, 2002] and GSN [Zhou and Troyanskaya, 2014]) on individual secondary structure labels in Table 2. Note that, in addition to the standard sequence feature and profile feature used by other methods (including ours) in this comparison, the CNF model [Wang *et al.*, 2011] was trained using three extra features, and the entire feature vector at an amino acid is 78-dimensional. Even though CNF was trained using more features, our model still achieves an overall higher Q8 accuracy. In terms of individual labels, our model achieves slightly lower accuracies (3% to 4% difference) on high frequency labels (H, E and L) and significantly higher accuracies on low frequency labels (T, S and B).

In order to further improve accuracy and robustness, we also compute an ensemble model by averaging 10 weak models trained independently on 10 randomly sampled training and validation subsets according to the bagging algorithm.

⁴<http://deeplearning.net/software/theano/>

⁵<http://keras.io/>

Table 2: A comparison of classification precisions of individual secondary structure labels on CB513. Note that CNF was trained with 78-dimensional features while GSN, DeepCNF and our model (DCRNN) were trained with the standard 42-dimensional features. The precision of a single model is calculated using the confusion matrix. The precision of an ensemble model is obtained from the voted confusion matrix. The frequencies of individual labels in the training set and their descriptions are presented in the last two columns. Boldface numbers indicate best performance.

	Single Model					Ensemble Model			Frequency	Description
	DCRNN	CNF	SSpro8	GSN	DeepCNF	DCRNN	CNF	SSpro8		
Q8	0.694	0.633	0.511	0.664	0.683	0.697	0.649	0.510		all types
H	0.836	0.887	0.752	0.831	0.849	0.832	0.875	0.752	0.309	α -helix
E	0.739	0.776	0.597	0.717	0.748	0.753	0.756	0.598	0.213	β -strand
L	0.573	0.608	0.499	0.518	0.571	0.573	0.601	0.500	0.211	irregular
T	0.549	0.418	0.327	0.496	0.530	0.559	0.487	0.330	0.118	β -turn
S	0.521	0.117	0.049	0.444	0.487	0.518	0.202	0.051	0.098	bend
G	0.432	0.049	0.007	0.450	0.490	0.429	0.207	0.006	0.037	3_{10} -helix
B	0.558	0.000	0.000	0.000	0.433	0.554	0.000	0.000	0.014	β -bridge
I	0.00	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	π -helix

Through the model averaging process, the Q8 accuracy can be improved to **69.7%** (4.8% higher than the previous best result from ensemble models), as shown in Table 2. In addition, the Q3 accuracy of our single model is **84.0%**, which is 1.7% higher than the previous state of the art (82.3%) [Wang *et al.*, 2016]. The mapping between 8-state labels and 3-state labels is as follows: H (8-state) is mapped to H (3-state), E (8-state) is mapped to E (3-state) and all other 8-state labels are mapped to C (3-state) according to [Wang *et al.*, 2016]. In addition, the p-value of the significance test, “our model outperforms other methods”, is 1.4×10^{-5} (< 0.001) using results from 10 different runs.

Training on Filtered CB6133 and Testing on CASP10 and CASP11

For further verifying the generalization capability of our model trained on the filtered CB6133 dataset, we have also evaluated it on more recent datasets, CASP10 and CASP11. We compare the Q8 accuracy of our model against SSpro [Magnan and Baldi, 2014], RaptorX-SS8 [Wang *et al.*, 2011] and DeepCNF. In addition, the Q3 accuracy of our model is also compared against SSpro, SPINE-X [Faraggi *et al.*, 2012], PSIPRED [Jones, 1999], RaptorX-SS8, JPRED [Drozdetskiy *et al.*, 2015] and DeepCNF. According to the results shown in Table 3, the Q8 accuracy of our model is **76.9%** over CASP10 (5.1% higher than previous best) and **73.1%** over CASP11 (0.8% higher than previous best). The Q3 accuracy of our model over CASP10 and CASP11 is respectively **87.8%** (3.4% higher than previous best) and **85.3%** (0.6% higher than previous best). Note that the Q8 accuracy and Q3 accuracy of our method in this test are obtained through a single model rather than an ensemble model for fair comparison.

3.4 Ablation Study

To discover the vital elements in the success of our proposed network, we conduct an ablation study by removing or replacing individual components in our network. Specifically, we have tested the performance of models without the feature embedding layer, multiscale CNNs, stacked BGRUs or backward RNN passes. In addition, we also tested a model

Table 3: Training on filtered CB6133 and testing on CASP10 and CASP11. Both Q8 and Q3 accuracies are reported for SSpro, SPINE-X, PSIPRED, JPRED, RaptorX-SS8, DeepCNF and our model (DCRNN). Boldface numbers indicate best performance.

Method	Q8(%)		Q3(%)	
	CASP10	CASP11	CASP10	CASP11
SSpro(without template)	64.9	65.6	78.5	77.6
SSpro(with template)	75.9	66.7	84.2	78.4
SPINE-X	-	-	80.7	79.3
PSIPRED	-	-	81.2	80.7
JPRED	-	-	81.6	80.4
RaptorX-SS8	64.8	65.1	78.9	79.1
DeepCNF	71.8	72.3	84.4	84.7
DCRNN	76.9	73.1	87.8	85.3

that does not feed local contexts to the last two fully connected layers, and another one where the BGRU layers are replaced with bidirectional simpleRNN layers without any gate structure to figure out the importance of the gate structure in presence of long-range dependencies. From the results on the CB513 dataset presented in Table 4, we find that those bidirectional GRU layers are the most effective component in our network as the performance drops to 66.9% when we run forward RNN passes only without backward RNN passes. Multiscale CNNs are also important as the performance drops to 68.1% without them. In addition, compared with bidirectional simpleRNN, the gate structure in our GRU layers is necessary for dealing with long-range dependencies widely existing in amino acid sequences. Stacked BGRU layers are also beneficial for enhancing global information circulation when compared with a single BGRU layer. Furthermore, directly feeding local contexts to the fully connected layers in addition to the global contexts is also essential for good performance, especially for the prediction of low-frequency secondary structure categories. Last but not the least, feature embedding, multi-task learning and bagging can all be applied to improve the accuracy and robustness of our method.

Table 4: An ablation study on CB513.

Model	Q8 Accuracy
Without feature embedding	68.9%
Without multiscale CNNs	68.1%
Without backward GRU passes	66.9%
Replacing GRU with simpleRNN	68.0%
Single BGRU layer	68.6%
Without integrated local&global contexts	68.8%
Without Multi-task learning	68.8%

4 Conclusions

To apply recent deep neural networks to protein secondary structure prediction, we have proposed an end-to-end model with multiscale CNNs and stacked bidirectional GRUs for extracting local and global contexts. Through the integrated local and global contexts, the previous state of the art in protein secondary structure prediction has been improved. By taking interactions between different protein properties into consideration, multi-task joint feature learning is exploited to further refine the performance. Because of the success of our proposed deep neural network on secondary structure prediction, such models combining local and global contexts can be potentially applied to other challenging structure prediction tasks in protein and computational biology.

Existing BGRUs are unable to deal with extremely long dependencies, especially low frequency long dependencies. More powerful architectures with an implicit attention mechanism, such as neural Turing machines [Graves *et al.*, 2014], may be suitable for solving this problem and further improving the prediction performance.

Acknowledgments

We wish to thank Sheng Wang and Jian Zhou for discussion and consultation about datasets, and the anonymous reviewers for their valuable comments. The first author is supported by Lee Shau Kee Postgraduate Fellowship from the University of Hong Kong.

References

- [Altschul *et al.*, 1997] Stephen F Altschul, Thomas L Madden, Alejandro A Schäffer, Jinghui Zhang, Zheng Zhang, Webb Miller, and David J Lipman. Gapped blast and psiblast: a new generation of protein database search programs. *Nucleic Acids Research*, 25(17):3390, 1997.
- [Asgari and Mofrad, 2015] Ehsaneddin Asgari and Mohammad RK Mofrad. Protvec: A continuous distributed representation of biological sequences. *arXiv preprint arXiv:1503.05140*, 2015.
- [Ashraf and Yaohang, 2014] Yaseen Ashraf and Li Yaohang. Context-based features enhance protein secondary structure prediction accuracy. *Journal of chemical information and modeling*, 54(3):992–1002, 2014.
- [Bastien *et al.*, 2012] Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, James Bergstra, Ian J. Goodfellow, Arnaud Bergeron, Nicolas Bouchard, and Yoshua Bengio. Theano: new features and speed improvements. 2012.
- [Bengio *et al.*, 1994] Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *Neural Networks, IEEE Transactions on*, 5(2):157–166, 1994.
- [Bergstra *et al.*, 2010] James Bergstra, Olivier Breuleux, Frédéric Bastien, Pascal Lamblin, Razvan Pascanu, Guillaume Desjardins, Joseph Turian, David Warde-Farley, and Yoshua Bengio. Theano: a CPU and GPU math expression compiler. 2010.
- [Cho *et al.*, 2014a] Kyunghyun Cho, Bart van Merriënboer Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder–decoder for statistical machine translation. page 1724C1734, 2014.
- [Cho *et al.*, 2014b] Kyunghyun Cho, Bart van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. On the properties of neural machine translation: Encoder–decoder approaches. *Syntax, Semantics and Structure in Statistical Translation*, page 103, 2014.
- [Chollet, 2015] Francois Chollet. Keras: Theano-based deep learning library. <https://github.com/fchollet/keras>, 2015.
- [Chou and Fasman, 1974] Peter Y Chou and Gerald D Fasman. Prediction of protein conformation. *Biochemistry*, 13(2):222–245, 1974.
- [Cuff and Barton, 1999] James A Cuff and Geoffrey J Barton. Evaluation and improvement of multiple sequence methods for protein secondary structure prediction. *Proteins: Structure, Function, and Bioinformatics*, 34(4):508–519, 1999.
- [Drozdetskiy *et al.*, 2015] Alexey Drozdetskiy, Christian Cole, James Procter, and Geoffrey J Barton. Jpred4: a protein secondary structure prediction server. *Nucleic acids research*, 43(W1):W389–W394, 2015.
- [Faraggi *et al.*, 2012] Eshel Faraggi, Tuo Zhang, Yuedong Yang, Lukasz Kurgan, and Yaoqi Zhou. Spine x: improving protein secondary structure prediction by multi-step learning coupled with prediction of solvent accessible surface area and backbone torsion angles. *Journal of computational chemistry*, 33(3):259–267, 2012.
- [Graves *et al.*, 2014] Alex Graves, Greg Wayne, and Ivo Danihelka. Neural turing machines. In *arXiv preprint arXiv:1410.5401*, 2014.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [Jones, 1999] David T Jones. Protein secondary structure prediction based on position-specific scoring matrices. *Journal of molecular biology*, 292(2):195–202, 1999.
- [Jozefowicz *et al.*, 2015] Rafal Jozefowicz, Wojciceh Zaremba, and Ilya Sutskever. An empirical exploration of recurrent network architectures. In *International Conference of Machine Learning*, Lille, France, 2015.
- [Kabsch and Sander, 1983] Wolfgang Kabsch and Christian Sander. Dictionary of protein secondary structure: pattern

- recognition of hydrogen-bonded and geometrical features. *Biopolymers*, 22(12):2577–2637, 1983.
- [Kingma and Ba, 2014] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [Kryshtafovych *et al.*, 2014] Andriy Kryshtafovych, Alessandro Barbato, Krzysztof Fidelis, Bohdan Monastyrskyy, Torsten Schwede, and Anna Tramontano. Assessment of the assessment: evaluation of the model quality estimates in casp10. *Proteins: Structure, Function, and Bioinformatics*, 82(S2):112–126, 2014.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [Magnan and Baldi, 2014] Christophe N Magnan and Pierre Baldi. Sspro/accpro 5: almost perfect prediction of protein secondary structure and relative solvent accessibility using profiles, machine learning and structural similarity. *Bioinformatics*, 30(18):2592–2597, 2014.
- [Mesnil *et al.*, 2015] Grégoire Mesnil, Yann Dauphin, Kaisheng Yao, Yoshua Bengio, Li Deng, Dilek Hakkani-Tur, Xiaodong He, Larry Heck, Gokhan Tur, Dong Yu, et al. Using recurrent neural networks for slot filling in spoken language understanding. *Audio, Speech, and Language Processing, IEEE/ACM Transactions on*, 23(3):530–539, 2015.
- [Moult *et al.*, 2014] John Moult, Krzysztof Fidelis, Andriy Kryshtafovych, Torsten Schwede, and Anna Tramontano. Critical assessment of methods of protein structure prediction (casp)-round x. *Proteins: Structure, Function, and Bioinformatics*, 82(S2):1–6, 2014.
- [Nair and Hinton, 2010] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pages 807–814, 2010.
- [Noble *et al.*, 2004] Martin EM Noble, Endicott Jane A, and Johnson Louise N. Protein kinase inhibitors: insights into drug design from structure. *Science*, 303(5665):1800–1805, 2004.
- [Pollastri *et al.*, 2002] Gianluca Pollastri, Darisz Przybylski, Burkhard Rost, and Pierre Baldi. Improving the prediction of protein secondary structure in three and eight classes using recurrent neural networks and profiles. *Proteins: Structure, Function, and Bioinformatics*, 47(2):228–235, 2002.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, pages 91–99, 2015.
- [Rost and Sander, 1993] Burkhard Rost and Chris Sander. Prediction of protein secondary structure at better than 70% accuracy. *Journal of molecular biology*, 232(2):584–599, 1993.
- [Sepideh *et al.*, 2010] Babaei Sepideh, Geranmayeh Amir, and Seyyedsalehi Seyyed Ali. Protein secondary structure prediction using modular reciprocal bidirectional recurrent neural networks. *Computer methods and programs in biomedicine*, 100(3):237–247, 2010.
- [Simpson and Morgan, 1983] Richard J Simpson and Francis J Morgan. Complete amino acid sequence of emben goose (*anser anser*) egg-white lysozyme. *Biochimica et Biophysica Acta (BBA)-Protein Structure and Molecular Enzymology*, 744(3):349–351, 1983.
- [Srivastava *et al.*, 2014] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.
- [Sujun and Zhirong, 2001] Hua Sujun and Sun Zhirong. A novel method of protein secondary structure prediction with high segment overlap measure: support vector machine approach. *Journal of molecular biology*, 308(2):397–407, 2001.
- [Wang and Dunbrack, 2003] Guoli Wang and Roland L Dunbrack. Pisces: a protein sequence culling server. *Bioinformatics*, 19(12):1589–1591, 2003.
- [Wang *et al.*, 2011] Zhiyong Wang, Feng Zhao, Jian Peng, and Jinbo Xu. Protein 8-class secondary structure prediction using conditional neural fields. *Proteomics*, 11(19):3786–3792, 2011.
- [Wang *et al.*, 2016] Sheng Wang, Jian Peng, Jianzhu Ma, and Jinbo Xu. Protein secondary structure prediction using deep convolutional neural fields. *Scientific reports*, 6, 2016.
- [Yaseen and Li, 2014] Ashraf Yaseen and Yaohang Li. Template-based c8-scorpion: a protein 8-state secondary structure prediction method using structural information and context-based features. *BMC bioinformatics*, 15(Suppl 8):S3, 2014.
- [Yih *et al.*, 2011] Wen-tau Yih, Kristina Toutanova, John C Platt, and Christopher Meek. Learning discriminative projections for text similarity measures. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 247–256. Association for Computational Linguistics, 2011.
- [Zhang *et al.*, 2015] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in Neural Information Processing Systems*, pages 649–657, 2015.
- [Zhou and Troyanskaya, 2014] Jian Zhou and Olga Troyanskaya. Deep supervised and convolutional generative stochastic network for protein secondary structure prediction. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 745–753, 2014.