

PREDICTING ANTICANCER PEPTIDES AND PROTEIN FUNCTION WITH
DEEP LEARNING

by

Nathaniel Patrick Lane

A thesis submitted in partial fulfillment
of the requirements for the degree

of

Master of Science

in

Computer Science

MONTANA STATE UNIVERSITY
Bozeman, Montana

April 2020

©COPYRIGHT

by

Nathaniel Patrick Lane

2020

All Rights Reserved

ACKNOWLEDGEMENTS

This work was made possible by the publicly available datasets by Balachandran et al. [19], Yi et al. [28], and Meher et al. [20]. I would like to thank Indika Kahanda for advising me through this project. I would also like to thank Montana State University's Gianforte School of Computing for providing the resources to complete this project, as well as the Bioinformatics group for their support.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. PREDICTING ANTICANCER PEPTIDES	3
Introduction	3
Related Work	5
<i>ACP-DL</i>	6
Methods.....	7
Datasets	7
<i>DeepACPpred</i> Model.....	11
Transfer Learning	15
Experimental Setup.....	15
Results.....	17
Sequence length distributions	17
Amino Acid compositions	17
Comparison to <i>ACP-DL</i>	18
Comparison to <i>MLACP</i>	19
Impact of Transfer Learning.....	21
Comparing the hybrid model to individual RNNs and CNNs models	22
Embeddings	23
Random Seeds.....	24
Conclusions and Future Work.....	26
3. PREDICTING PROTEIN FUNCTION.....	28
Introduction	28
Related Work	29
Methods.....	30
Datasets	30
Adapting <i>DeepACPpred</i>	31
Training.....	31
Results and Discussion	34
Conclusion and Future Work	37
4. CONCLUSION AND FUTURE WORK.....	39
REFERENCES CITED.....	41

TABLE OF CONTENTS – CONTINUED

APPENDICES	45
APPENDIX A : PROTEINS.....	46
Amino Acids	47
Protein Folding.....	48
APPENDIX B : NEURAL NETWORKS.....	52
Dense Layers	53
Activation	54
Loss	55
Gradient Descent	55
Embeddings	56
Word2vec	56
Recurrent Neural Networks	57
Long Short-Term Memory Layers	58
Convolutional Networks	60
Vanishing Gradients and Overfitting	60
Dropout	61
Regularization	61
Batch Normalization	63
Performance Metrics	63
Cross Validation	64
APPENDIX C : PROTEIN LENGTHS FOR ALL DATASETS.....	65

LIST OF TABLES

Table	Page
2.1 An analysis of the make-up of each dataset.....	8
2.2 Comparison of <i>DeepACPpred</i> against <i>ACP-DL</i> using the ACP240 dataset with 5-Fold Cross-Validation. The best values are indicated in bold text.....	18
2.3 Comparison of <i>DeepACPpred</i> against <i>ACP-DL</i> using the ACP740 dataset with 5-Fold Cross-Validation. The best values for are indicated in bold text.....	19
2.4 Comparison of <i>DeepACPpred</i> against <i>ACP-DL</i> using the Tyagi dataset with 5-Fold Cross-Validation. The best values are indicated in bold text.....	19
2.5 MCC scores for each fold of <i>ACP-DL</i> 's experimentation on the ACP740 dataset [28]	20
2.6 Comparison of <i>DeepACPpred</i> against <i>MLACP</i> model using the Tyagi dataset and 10-Fold Cross Validation. The best values are indicated in bold text.	20
2.7 Comparison of <i>DeepACPpred</i> against <i>MLACP</i> and other models using the Tyagi training set and HC testing set. The best values are indicated in bold text.	21
2.8 Comparison of <i>DeepACPpred</i> against <i>MLACP</i> and other models using the Tyagi training set and LEE testing set. The best values are indicated in bold text.	22
2.9 Comparing base RNN, base CNN, and combined CNN + RNN models using the Tyagi dataset and 10-Fold Cross Validation. The best values are indicated in bold.	23
3.1 F_{max} scores of this model when testing on CAFA 3 result data and performing 5-fold cross validation with training data.	36
A.1 Amino acids with their abbreviations	49
B.1 Performance metrics.....	63

LIST OF FIGURES

Figure	Page
2.1 A comparison of the sequence lengths in each dataset.....	9
2.2 The fraction of each dataset that is made up by each amino acid.....	10
2.3 Structure of the final model.	13
2.4 An analysis of the make-up of each dataset.....	14
2.5 Evaluation of the effects of transfer learning using the Tyagi dataset and 10-Fold Cross Validation	22
2.6 A comparison of the MCCs obtained with different random seeds. The blue bars represent the results from 10-fold cross validation on the Tyagi dataset, while the red is for testing on the HC dataset.	25
3.1 A histogram comparison of the sequence lengths in a selection of unfiltered datasets.....	32
3.2 A histogram comparison of the sequence lengths in a selection of filtered datasets.	33
A.1 Basic amino acid structure.....	47
A.2 Condensation reaction.	48
A.3 Secondary structure of a protein.....	50
A.4 Tertiary structure of a protein.....	50
A.5 Quaternary structure of a protein.....	51
B.1 A simple example of a neural network	53
B.2 Plot of the sigmoid function.....	54
B.3 Representation of skip-gram.....	57
B.4 Representation of CBOW	57
B.5 Diagram showing how an RNN works.....	58
B.6 Diagram showing how an LSTM works.....	59

LIST OF FIGURES – CONTINUED

Figure	Page
B.7 Graph showing the difference between a complex function and a simple one.	62
B.8 Graphic showing the process of cross validation.....	64
C.1 A histogram comparison of the sequence lengths in each unfiltered dataset.....	66
C.2 A histogram comparison of the sequence lengths in each filtered dataset.	69

ABSTRACT

Anticancer peptides (ACPs) are a promising alternative to traditional chemotherapy. To aid wet-lab and clinical research, there is a growing interest in using machine learning techniques to help identify good ACP candidates computationally. In this work, we develop *DeepACPpred*, a novel deep learning model for predicting ACPs using their amino acid sequences. Using several gold-standard ACP datasets, we demonstrate that *DeepACPpred* is highly effective compared to state-of-the-art ACP prediction models. Furthermore, we adapt the above neural network model for predicting protein function and report our experience with participating in a community-wide large-scale assessment of protein functional annotation tools.

INTRODUCTION

Proteins have a broad number of functions within the body. These can range from facilitating movement as muscle fibers to speeding up chemical reactions as enzymes. Proteins are formed by chaining together a number of smaller molecules classified as amino acids. Each amino acid has a different chemical makeup and therefore different properties. One of the main determining factors for the structure and function of a protein is this amino acids sequence. Relevant biological background is provided in Appendix A.

Lately, there has been an interest in using computers to predict these functions from the chemical structure of the proteins. Many previous works have used properties of the whole protein to accomplish this. This work presents a technique for doing this by analyzing protein's amino acid sequence. Particular emphasis is placed on predicting whether a given sequence of amino acids has cancer-fighting properties using deep neural networks. Relevant computational background on neural networks is provided in Appendix B.

The core part of the work presented in this work is laid out in chapter 2. That chapter outlines *DeepACPPred*, a system based on recurrent neural networks that predicts whether or not a given sequence of amino acids represents an anticancer peptide (a short chain of amino acids that could be thought of as a small protein). Chapter 2 is an extended version of our paper recently accepted for the 14th International Conference on Practical Applications of Computational Biology and Bioinformatics (PACBB).

The fundamental problem solved by *DeepACPPred* is predicting if an amino

acid sequence has certain properties. The input does not necessarily have to be anticancer peptides; it could be any amino acid sequence which has some desired property. As such, in chapter 3, we adapt the system developed in chapter 2 to predicting a protein's Gene Ontology annotations. The Gene Ontology is a hierarchical system which describes the qualities and functions which a protein might have. This was done as participation in a competition called the Critical Assessment of protein Function Annotation algorithms (CAFA).

Ultimately, this work aims to show that deep neural networks are a promising technique for predicting the properties of peptides and proteins using their amino acid sequences. Such computational models have significant implications for biologists and clinicians for expediting wet-lab and clinical experiments, which are known to be extremely resource consuming.

PREDICTING ANTICANCER PEPTIDES

This chapter is an extended version of the paper titled “DeepACPpred: A Novel Hybrid CNN-RNN Architecture for Predicting Anti-Cancer Peptides” recently accepted to the 14th International Conference on Practical Applications of Computational Biology & Bioinformatics.

Introduction

Chemotherapy, one of the primary treatments for cancer, often has side effects that are debilitating for the patient. Some examples of problems that can arise include vomiting, hair loss, and fatigue [4]. Additionally, due to the fact that cancer cells reproduce at an unregulated pace, tumors will often develop resistances to chemotherapeutic drugs (In fact, it is estimated that this is the cause of 90% of treatment failures) [18]. These facts, taken together with the existence of many different kinds of cancer, show why there is a need for the development of many kinds of anticancer drugs.

Recent research shows that anticancer peptides (ACPs) may offer a promising alternative to chemotherapy [11, 12]. These peptides are typically sequences of 5-30 amino acids that exhibit physicochemical properties which help them target cancer cells [6]. To aid wet-lab and clinical researchers, there has been a recent interest in developing machine learning algorithms that can help identify good ACP candidates [3, 19, 22].

Most of these predictive models use features that are generated from the physicochemical properties of the amino acids that comprise the ACP. We hypothesize that a model that leverages the actual sequences of amino acids may outperform these other models. Our reasoning is based on the fact that the sequence determines the

folding structure, and therefore, the function of the peptide; hence, the sequence itself would be more informative than features generated using the sequence properties.

Therefore, we propose to use deep learning for developing a model for predicting ACPs. Deep learning eliminates the need for tedious feature engineering and the features are learned by the model itself based on the “raw” sequence input. Our proposed method of ACP prediction, *DeepACPred*, is specifically based on a hybrid model composed of Recurrent Neural Networks (RNNs) and Convolutional Neural Networks (CNNs). The intuition for this design is that a CNN may be able to extract information from interactions among nearby amino acids, and that an RNN may be more capable of analyzing interactions among distant amino acids. Additionally, the CNN removes information regarding the specific amino acids involved in the convolution, so the RNN can ideally retain the data that is lost.

We compare the performance of the proposed model against the models from two state-of-the-art ACP prediction studies on several popular ACP gold-standard datasets. The first of these studies analyzed the performance of both Random Forest (RF) and Support Vector Machine (SVM) techniques, and, at the time of writing, appears to be the best model for ACP prediction [19]. The other study is the first (and to our knowledge at the time of writing, the only) RNN-based approach to predicting ACPs [28]. Our experimental results indicate that our proposed model is highly effective for ACP prediction in comparison to the models from the above two studies.

The main contributions presented by our work are (1) the introduction of a novel CNN-RNN hybrid architecture for ACP prediction, (2) the exploration of the ability of embeddings to contribute to this problem, (3) validation of the importance of transfer learning for this particular domain. The rest of this chapter is organized as follows. The Related Work section discusses previous work conducted on the problem

of ACP prediction. In the Methods section, we detail in-depth our approach to this problem. The Results section reports and analyzes the performance of our model and compares it to other state-of-the-art methods. Finally, in the Conclusions and Future Work section, we report on what work may be performed in the future to improve the techniques built in this work.

Related Work

Manavalan et al. [19] present an SVM model for predicting whether or not a given sequence of amino acids represents an ACP. Their proposed solution, *MLACP*, uses features gained from the physiochemical properties of the amino acids, as well as the raw amino acid composition. *MLACP* will be used as a baseline for evaluating the performance of our model, and we will be using the same datasets and metrics that they used. Since they compared their model to *AntiCP* [22] and *iACP* [3], both of which are SVM-based models, and provided metrics against them, we do the same when evaluating *DeepACPpred*.

The core of this problem is protein classification. Since there has not yet been a lot of research done on using RNNs to predict whether or not a given peptide has anticancer properties, it is important to identify previous research into using RNNs to classify peptides and proteins of other types. Hamid et al. [10] present such research. In this paper, the researchers essentially use the skip-gram Word2vec algorithm on trigrams within amino acid sequences to generate embeddings. These embedding vectors are then used as the input to the RNN. This method saw a modest improvement over other machine learning techniques.

ACP-DL

One very recent paper by Yi et al. [28] does directly address the application of RNNs to the prediction of ACPs. In their proposed method, *ACP-DL*, each amino acid in each sequence is converted into two vectors which are concatenated.

The first vector is generated via a Binary Profile Feature (BPF), wherein each of the twenty amino acids (A, C, D, E, F, G, H, I, K, L, M, N, P, Q, R, S, T, V, W, and Y) is represented with a one-hot vector. All elements within this vector are 0s except for a single 1 in the spot corresponding with the amino acid’s place in the alphabet. For instance, the amino acid A will be replaced with the vector (1,0,...,0), and the amino acid C will be replaced with the vector (0,1,...,0). Once this is done, the vectors in every window of 7 amino acids will be concatenated into a vector of length 140.

The second is a k -mer sparse matrix, which is a technique these researchers had previously developed [29]. Here, each amino acid is placed into one of seven groups based on its dipole moment and side chain volume. The groups are as follows:

- A, G, and V
- I, L, F, and P
- Y, M, T, and S
- H, N, Q, and W
- R and K
- D and E
- C

Each amino acid is replaced with the id of its respective group. Now, considering subsequences of k amino acids (a k -mer), there are 7^k such possible subsequences. Assuming the length of a peptide is L , then there are $L - k + 1$ k -mers within that peptide. Each of these k -mers is converted to a 7^k vector, which is all 0s except for a single 1 at the point in the vector corresponding with that particular k -mer (much like in the Binary Profile Feature technique outlined above). In this particular project, they used a k of 3. Finally, once all of these vectors have been generated, they are concatenated and fed into a deep LSTM.

To our knowledge, *ACP-DL* is the only other paper to directly use RNNs to predict whether or not an input string of amino acids represents an ACP; however, some other studies present similar work. For instance, in two studies, Grisoni et al. use RNNs [7] and counter-propagation artificial neural networks (CPANNs) [8] to generate ACPs which were then experimentally verified in a wet-lab setting. While these studies do present work similar to ours, the difference between peptide “generation” and “prediction” is not an insignificant one. A very recent study, *PTPD* by Wu et al., used Word2vec embeddings with a fully convolutional network to tackle this problem [27]. We attempted to integrate these techniques into *DeepACPPred*, but ultimately settled on the hybrid CNN-RNN model.

We will compare our results to the those of *ACP-DL* as well as *MLACP*.

Methods

Datasets

There are six datasets that are used for training and testing in this project: the Tyagi, LEE, and HC datasets from Manavalan et al. [19]; the ACP240 and ACP740 datasets from Yi et al. [28]; and the final key dataset is the anti-fungal peptide (AFP) dataset from Meher et al. [20]. Table 2.1 shows the balance of positive to negative


Dataset	# Positives	# Negatives	Percent Positive	Reference
Tyagi	187	399	32%	[19]
LEE	422	422	50%	[19]
HC	126	205	38%	[19]
ACP240	129	111	54%	[28]
ACP740	376	364	51%	[28]
AFP	1497	1393	52%	[20]

Table 2.1: An analysis of the make-up of each dataset.


examples in each dataset. The lengths of the peptide sequences in each dataset are shown in Figure 2.1, and the amino acid compositions in Figure 2.2.

The Tyagi dataset was created for MLACP by combining two previously created datasets: Balanced 1 (B1) and Balanced 2 (B2). The paper originating these datasets, by E-Kobon et al., analyzed peptides found in the mucus of a snail (*Achatina fulica*) to see if they showed cytotoxicity against breast cancer cells [5]. The Hajisharifi-Chen (HC) dataset was previously used in papers by both Hajisharifi and Chen (the author of *iACP*). Hajisharifi et al. generated the set by getting a list of antimicrobial peptides from the antimicrobial peptides database created by Wang et al. [24] and by searching articles. The justification for this is that many antimicrobial peptides are also anticancer peptides [6]. The negative samples were gathered by selecting random samples from UniProtKB [9]. The LEE dataset was created for *MLACP* by merging a dataset of ACPs reported by iACP [3] with experimentally verified ones from CancerPPD [23] and APD3 [25]. The non-ACPs of this dataset were random peptides from Swiss-prot [19].


The ACP740 and ACP240 datasets were created by merging results from




(a) Lengths of peptides in Tyagi




(b) Lengths of peptides in LEE




(c) Lengths of peptides in HC



(d) Lengths of peptides in ACP240




(e) Lengths of peptides in ACP740




(f) Lengths of peptides in AFP


Figure 2.1: A comparison of the sequence lengths in each dataset.




(a) Amino acid composition of Tyagi




(b) Amino acid composition of LEE




(c) Amino acid composition of HC



(d) Amino acid composition of ACP240



(e) Amino acid composition of ACP740



(f) Amino acid composition of AFP

Figure 2.2: The fraction of each dataset that is made up by each amino acid.

iACP [3] and the work of Wei et al. [26] with no overlap between the resulting two sets [28]. The AFP dataset was created by merging results downloaded from LAMP [30], an AMP database, and APD3 [25] with predictions from LAMP[30], AntiBP2 [17], and AVPpred [21].

DeepACPpred Model

Over the course of this project, there have been several iterations of the neural network structure. The final structure takes the amino acid sequences as input and then feeds that into two different paths. Each path begins with an embeddings layer and ends with a BiLSTM; the difference is that one path has a 1D convolutional layer in the middle. Conceptually, the addition of this layer should allow the network to account for interactions among amino acids which are spatially close. This comes at the cost of potentially removing information that could be useful for analyzing interactions among amino acids that are distant, so the solution is to have the one path which has the convolution and one that does not; the latter may be able to analyze the whole sequence without the removal of data via convolution. When we experimented with the CNN-MaxPool structure presented in *PTPD* [27] for this branch, we found that it offered no significant improvements to the CNN-RNN structure shown here. Figure 2.3 shows an outline of the neural network structure.

The first layer along both paths is an embeddings layer, which converts each amino acid to a unique vector. This is done by creating a matrix whose width is the size of the vocabulary (in this case, 22; 20 for the amino acids, 1 for a sequence terminator, and 1 for padding) and whose height is the size of the output of the layer. Before being fed into this layer, each amino acid must be converted into an integer in the range (0, 22]. The layer uses this number to create a one-hot vector, which is a vector whose values are all 0s, except for a single 1. The length of the vector is equal

to the number of elements in the vocabulary (in this case, 22), and the 1 is found at the index whose value is the amino acid’s unique integer. The matrix and vector are then multiplied together. When given an amino acid whose number is n , the n th column of the matrix will be the output of this operation. Thus is each amino acid converted to a vector.

We capped off each sequence with a unique integer that was not assigned to an amino acid. This was used to signal to the network that the sequence had terminated. Further, each one of the amino acid sequences must have the same length to be used in an RNN, so we padded them with another unique integer. Ideally, the network would learn that this identifier was not a part of the amino acid sequence.

The values in the embeddings layer matrix can be fully-trainable (i.e., ones that start random and the network can adjust during training) or pre-trained and fixed. We experimented with both methods in this project. The pre-trained embeddings were produced by decomposing each amino acid sequence of the experimentally-verified human peptides on UniProt into a series of trigrams. The Word2vec skip-gram algorithm was applied to these trigrams, producing embedding vectors. See appendix B for a detailed explanation of this algorithm.

The model was implemented with the Keras¹ library. This network is compiled with the *Adam* optimizer and a Binary Crossentropy loss function. To help prevent overfitting, regularization is used in addition to *dropout*. While batch normalization typically produces better performance than either of these techniques [13], we found that in this case, it produced inferior results. Because of the large number of hyperparameters, this model is difficult to optimize by hand. Therefore, we employed the use of the library *hyperas*² to automatically optimize the hyperparameters.

¹<https://keras.io/>

²<https://maxpumperla.com/hyperas/>




Figure 2.3: Structure of the final model.

Layer	Output Shape	Number of Params	Connected to
Input	(131)	0	-
Embeddings (CNN)	(131, 55)	1155	Input
Dropout (CNN)	(131, 55)	0	Embeddings (CNN)
Conv1D (CNN)	(113, 83)	86,818	Dropout (CNN)
Dropout 2 (CNN)	(113, 83)	0	Conv1D (CNN)
BiLSTM (CNN)	(250)	209,000	Dropout 2 (CNN)
Embeddings (RNN)	(131, 144)	3024	Input
Dropout (RNN)	(131, 144)	0	Embeddings (RNN)
BiLSTM (RNN)	(212)	212,848	Dropout (RNN)
Concatenate	(462)	0	[BiLSTM (CNN), BiLSTM (RNN)]
Dropout	(462)	0	Concatenate
Dense	(100)	46,300	Dropout
Dropout 2	(100)	0	Dense
Dense 2	(1)	101	Dropout 2

Figure 2.4: An analysis of the make-up of each dataset.

Transfer Learning

None of the provided ACP datasets are particularly large; the biggest is LEE with 844 amino acids sequences. This poses a problem, as neural networks tend to perform best when the training set has a high number of samples to draw from.

Our solution was to implement transfer learning. This is a process wherein the neural network is trained at first with a dataset that, while not directly pertaining to the problem at hand, is related to the data being tested. Once trained, the model is saved, and at the beginning of each iteration of cross validation, the model is loaded and can be trained on ACP data.

In this case, we used the AFP dataset presented in [20]. This was chosen because some anti-microbial peptides have anticancer properties [6], and this dataset has 1,496 positive entries and 1,384 negative entries, for a total number of 2,880. This is a much greater amount of data than is present in any of the ACP datasets.

Once the neural network has been trained, the weights are saved and transferred to a neural network with the same overall structure. This new network is then trained and tested on ACP data.

To show that transfer learning does indeed help, we will compare the results with and without it. These results will be obtained by doing 10-fold cross validation on the Tyagi dataset.

Experimental Setup

Because we are comparing the quality of our model to those proposed by a couple of other studies, we use a number of different evaluation schemes to mimic theirs as closely as possible. First, to compare with *MLACP* [19], we will optimize our model’s hyperparameters based on 10-fold cross-validation performed on the Tyagi dataset. For testing, the model is trained with the Tyagi dataset and evaluated against the

LEE and HC datasets. We will then compare our numbers directly with theirs. Since their metrics were MCC, accuracy, sensitivity, and specificity, we will use those as well. Ideally, we would also have been able to have used *MLACP* with other datasets for our comparison; however, we were never afforded access to the code. Nevertheless, throughout the course of this work, we have been able to train and test *DeepACPPred* on many datasets, so we are confident of its generalizability.

When comparing with *ACP-DL*, we will use 5-fold cross validation on the ACP240 and ACP740 datasets, just as they did. Their performance metrics were accuracy, sensitivity, specificity, precision, and MCC, so we will evaluate our performance similarly. Even though they do not present F1 scores in their study, they provide sensitivity and precision, from which F1 can be calculated, so we provide those as well. Additionally, since we had access to the source code of this tool, we were able to use it on other datasets. Specifically, we evaluated its performance with the Tyagi dataset.

As part of our experimentation, we will show that the hybrid model is superior to the RNN alone and the CNN alone. This will help support our intuition that the combination of the two is a better model than either one alone. We will do this by optimizing an RNN, a CNN, and the hybrid via 10-fold cross validation on the Tyagi dataset, and presenting the results. We will also use a similar method to evaluate the effect of transfer learning.

Further, we explore the question of how much variance is in our model, i.e., to what degree changing the random seed will affect the outcomes of the experiments. To answer this, we will perform two tests. The first will do 10-fold cross validation on the Tyagi dataset using 10 random seeds. The second will train on the Tyagi dataset and test on the HC dataset with the same 10 seeds. The first test will show how much variance there is during validation, while the second will show how much there

is during testing.

Results

Sequence length distributions

When comparing our results to those of *MLACP*, we train on the Tyagi dataset and test on both the HC and LEE datasets. Because we are training and testing on separate datasets, it is worth comparing the properties of the datasets for gaining intuition.

The lengths of the sequences in Tyagi and HC are similar for both positives and negatives (see Figure 2.1). That said, the distribution of lengths in the HC dataset is slightly narrower than in Tyagi. This may be due to the fact that HC, with 331 samples, is smaller than Tyagi, which has 586 samples. Regardless, the distribution of amino acid sequence lengths is the most similar between Tyagi and HC.

Importantly, there is a large disparity in the lengths of amino acid sequences seen in LEE vs Tyagi. In Tyagi, the distribution of positives appears to be bimodal with a preference for shorter sequences, while the distribution for the negatives appears to be approximately a normal distribution skewed right. In the LEE dataset, the positives are more significantly skewed right, with the distribution favoring shorter sequences even more than in Tyagi. The negatives in this dataset appear to follow an almost perfectly bimodal distribution that is quite broad. This is significantly different from the Tyagi dataset.

Amino Acid compositions

When looking at the amino acid distributions of the datasets (Figure 2.2), it again becomes apparent that Tyagi is more similar to HC than it is to LEE. Between Tyagi and HC, the biggest difference appears to come from the amino acids

I (Isoleucine) and *K* (Lysine). Both are less represented in HC’s negative samples than in Tyagi’s.

The amino acid distribution between LEE and Tyagi is much more significantly different. While the distribution for negative samples appears to be similar, the positive samples seem to have a much higher concentration of *K* (Lysine) and *L* (Leucine) than in Tyagi. There is also a much lower amount of *C* (Cysteine) in LEE than in Tyagi. This disparity will likely have a negative impact on any machine learning model.

Comparison to *ACP-DL*

Tables 2.2, 2.3, and 2.4 depict the performance of *DeepACPpred* against *ACP-DL* on the ACP240, ACP740, and Tyagi datasets respectively. We used 5-fold cross evaluation to evaluate the model to maintain consistency with their study. *DeepACPpred* outperformed *ACP-DL* in almost every performance metric except for specificity on the ACP240 dataset.

Algorithm	MCC	Accuracy	Sensitivity	Specificity	Precision	F1
ACP-DL	0.714	0.854	0.846	0.899	0.803	0.824
DeepACPpred	0.716	0.858	0.880	0.840	0.862	0.869

Table 2.2: Comparison of *DeepACPpred* against *ACP-DL* using the ACP240 dataset with 5-Fold Cross-Validation. The best values are indicated in bold text.

The authors of *ACP-DL* provided the results of each fold in their analysis [28], shown here for convenience in table 2.5. Using the MCC scores for the ACP740 dataset, we calculated the significance of our results using the Student’s *t*-test with 4 degrees of freedom (the number of folds minus one). The result is a *p*-value of 0.027.

Algorithm	MCC	Accuracy	Sensitivity	Specificity	Precision	F1
ACP-DL	0.631	0.815	0.826	0.806	0.824	0.825
DeepACPpred	0.706	0.850	0.853	0.850	0.856	0.850

Table 2.3: Comparison of *DeepACPpred* against *ACP-DL* using the ACP740 dataset with 5-Fold Cross-Validation. The best values for are indicated in bold text.

Algorithm	MCC	Accuracy	Sensitivity	Specificity	Precision	F1
ACP-DL	0.504	0.786	0.626	0.862	0.691	0.657
DeepACPpred	0.608	0.832	0.716	0.886	0.743	0.724

Table 2.4: Comparison of *DeepACPpred* against *ACP-DL* using the Tyagi dataset with 5-Fold Cross-Validation. The best values are indicated in bold text.

This is less than the common threshold of 0.05, which can be interpreted as meaning that if our results were not actually significant, then our results could be expected to be seen only 5% of the time. Thus, we can say with 95% accuracy that our results were indeed significantly better than theirs when experimenting with the ACP740 dataset.

Comparison to *MLACP*

Table 2.6 shows the performance of our model against the RFACP and SVMACP models presented in the *MLACP* study [19]. Each model was evaluated using 10-fold cross validation on the Tyagi dataset. *DeepACPpred* is able to outperform their models on Sensitivity.

The *MLACP* study outlined an additional testing procedure, in which their models were trained on the Tyagi dataset and tested on the HC and LEE datasets.

Fold number	MCC
1	0.5958
2	0.6639
3	0.6222
4	0.7163
5	0.5547

Table 2.5: MCC scores for each fold of *ACP-DL*’s experimentation on the ACP740 dataset [28]

Algorithm	MCC	Accuracy	Sensitivity	Specificity
RFACP	0.698	0.872	0.722	0.942
SVMACP	0.697	0.872	0.706	0.95
DeepACPpred	0.655	0.854	0.729	0.913

Table 2.6: Comparison of *DeepACPpred* against *MLACP* model using the Tyagi dataset and 10-Fold Cross Validation. The best values are indicated in bold text.

With this approach, the authors were able to compare their results against previous models presented in other studies [3, 22]. Tables 2.7 and 2.8 compare the performance of *DeepACPpred* against all models presented in the *MLACP* paper. Our model outperformed all models that had been described prior to their work, while outperforming their RF model in sensitivity when tested with the HC dataset, and being competitive in terms of sensitivity compared to their models on the LEE dataset.

HC tends to have significantly better performance than LEE. This is likely because, as explained above, it is more similar to the training set than LEE is. Of

Algorithm	MCC	Accuracy	Sensitivity	Specificity
RFACP	0.885	0.946	0.889	0.981
DeepACPpred	0.843	0.924	0.944	0.912
SVMACP	0.750	0.882	0.841	0.907
AntiCP (Model 2) [22]	0.719	0.869	0.813	0.902
AntiCP (Model 1) [22]	0.062	0.402	0.976	0.049

Table 2.7: Comparison of *DeepACPpred* against *MLACP* and other models using the Tyagi training set and HC testing set. The best values are indicated in bold text.

particular note, the sequence length differences between LEE and Tyagi may also explain why *DeepACPpred* has a greater dip in performance than other methods when evaluating on LEE. The other models are based on SVM and RF machine learning algorithms, and while the length of the sequences may alter some of the input features for these methods, *DeepACPpred*, being based on RNNs, is perhaps more sensitive to sequence length than the other models.

While *ACP-DL* provided the results for each fold, from which we could show the significance of our results, *MLACP* offered no such data. Thus, we were unable to determine if *MALACP* had significantly better performance than *DeepACPpred*.

Impact of Transfer Learning

Figure 2.5 shows the results of performing 10-fold cross validation on the Tyagi dataset with and without transfer learning from the AFP dataset. The application of transfer learning raised the MCC by 21.52%, accuracy by 7.69%, sensitivity by 2.54%, and specificity by 9.74%. This confirms our hypothesis that using anti-fungal peptides as a transfer learning set greatly impacts the success of the model.

Algorithm	MCC	Accuracy	Sensitivity	Specificity
RFACP	0.674	0.827	0.706	0.948
SVMACP	0.630	0.814	0.775	0.853
DeepACPpred	0.578	0.789	0.774	0.803
AntiCP (Model 2) [22]	0.505	0.752	0.744	0.761
iACP [3]	0.412	0.706	0.697	0.716
AntiCP (Model 1) [22]	0.096	0.527	0.938	0.116

Table 2.8: Comparison of *DeepACPpred* against *MLACP* and other models using the Tyagi training set and LEE testing set. The best values are indicated in bold text.

	MCC	Accuracy	Sensitivity	Specificity
Without Transfer	0.539	0.793	0.711	0.832
With Transfer	0.655	0.854	0.729	0.913

Figure 2.5: Evaluation of the effects of transfer learning using the Tyagi dataset and 10-Fold Cross Validation

Comparing the hybrid model to individual RNNs and CNNs models

Table 2.9 compares the results of 10-fold cross validation on the Tyagi dataset when using RNN, CNN and hybrid CNN + RNN models. This demonstrates that the combined model performs better than either network alone and supports our hypothesis that the networks train on separate aspects of the amino acid sequences.

Model	MCC	Accuracy	Sensitivity	Specificity
RNN	0.642	0.844	0.783	0.871
CNN	0.575	0.816	0.679	0.883
RNN + CNN	0.655	0.854	0.729	0.913

Table 2.9: Comparing base RNN, base CNN, and combined CNN + RNN models using the Tyagi dataset and 10-Fold Cross Validation. The best values are indicated in bold.

Embeddings

Along both paths (Figure 2.3), the amino acid sequence is first fed into an embeddings layer. Initially, we believed that fully trainable embeddings would too greatly increase the number of parameters in the model and lead to overfitting, so we experimented with both pre-trained and untrained embeddings. Word2vec embeddings are powerful tools, especially in language processing, so we believed that they may bolster the performance of our model. Contrary to our initial hypothesis, we discovered that initially random, fully-trainable embeddings performed best.

One hypothesis as to why the pre-trained embeddings did not work as well is that the set from which they were generated was too small. The pre-trained embeddings were produced by decomposing each amino acid sequence of the experimentally-verified human peptides on UniProt into a series of trigrams. The Word2vec skip-gram algorithm was applied to these trigrams, producing embedding vectors. The set of experimentally-verified human peptides only has 81 entries. We did also experiment with experimentally-verified peptides from all species, but these embeddings did not perform much better.

The second hypothesis is that the network is adjusting the trainable embeddings

based on different criteria from the Word2vec algorithm. It is famously difficult to determine exactly how a neural network is behaving. Nevertheless, the Word2vec embeddings are designed to determine how an amino acid interacts with its neighbors. It does not reflect how it interacts with the whole sequence. The vectors produced by an embeddings layer in a neural network, however, are more suited to the behavior of that network. Since the LSTM carries information across the whole sequence, it will be better for the embeddings vectors to consider the amino acid’s role in that sequence, and that is precisely how the values of the vectors will be determined by an embeddings layer.

We believe that this is the more likely of the two hypotheses based on the outcomes of some of our experiments. One reason for this is that changing the set that the embeddings were training on had a minimal effect on the quality of the neural network. The other reason is that the generated embeddings likely encode for information that the neural network is not looking at in the fully-trainable version. It has the freedom to explore other elements of the data that are not represented in the pre-trained embeddings.

Random Seeds

To reveal how much variance is in our model, we performed two tests. The first was 10-fold cross validation on the Tyagi dataset 10 times with the seeds 0-9 inclusive. The second trained on the Tyagi dataset and tested on the HC dataset 10 times with the same seeds. Figure 2.6 compares the resulting MCC scores for each test.

Interestingly, we found that while for most seeds, testing on the HC dataset performs about as well as the results presented above in table 2.7 (which reports results obtained using a single randomly generated seed), a few seeds give performance that exceeds even that of *MLACP*. The averaged MCC (across all seeds) of our model




Figure 2.6: A comparison of the MCCs obtained with different random seeds. The blue bars represent the results from 10-fold cross validation on the Tyagi dataset, while the red is for testing on the HC dataset.

is 0.887, which actually exceeds the value (i.e. 0.843) reported in table 2.7, and even slightly exceeds the performance of *MLACP*.

However, note that neither *MLACP* nor *ACP-DL* report results from repeated tests or analyze how different random seeds might affect the performance of their models. Our experimental evidence, however, reveals that such variance studies may be helpful in determining the true strength of each model. We suggest that the community could benefit from a study comparing the variance among many ACP-predicting models. We also acknowledge that the impact of random seeds on the performance may be minimal on traditional machine learning models such as SVMs compared to deep learning models.

Conclusions and Future Work

In this work, we introduce a novel CNN-RNN hybrid architecture for ACP prediction, explore the effectiveness of various embeddings, and demonstrate the importance of transfer learning for this particular domain. As far as we know, *DeepACPpred* is only the second model to use deep learning to predict ACPs. While it offers a great improvement over the previous RNN model, it is also competitive with the other current state-of-the-art models. In this section, we will highlight some of our findings in the hope that further refinement of the work presented here may offer a great improvement over the current best models.

Firstly, in this project, only the amino acid sequence itself was used as input to the neural network. It may be that adding additional features to the inputs may improve the results. One particular set of features that we have in mind is the annotations from the Gene Ontology (GO) for each amino acid sequence. Second, it is possible that other transfer learning sets could have more success than the anti-fungal peptides. This was really the only set we experimented with, so additional

tests may reveal that there is a better one to use.

Additionally, the current pool of available ACP data is rather small. Because of this, to prevent overfitting, our model had to have relatively few parameters. As the amount of available ACP data grows, we hope that we will be able to similarly grow the number of parameters in *DeepACPred* as well, ultimately improving performance. Further, we believe that current ACP data could benefit from a verification study by biologists. Alongside this, we believe that creating a database of verified non-ACPs could be helpful. Most datasets used in this project generated their negatives not from actual research data, but by randomly selecting amino acid sequences from UniProtKB. While it is true that most peptides will not be ACPs, it cannot be guaranteed that all of the elements used as negative examples are actually non-ACPs. Overall, we believe that better training and testing sets would be highly beneficial.

Finally, our results indicated a fair amount of variance in the outcomes of our model. Since none of the other models that were analyzed in this project reported variance across multiple runs, we propose that the community could benefit from a study analyzing how much variance these models have. In doing so, we could gain a better understanding of the true performance of these models.

For *DeepACPred* to be useful to the medical research community, a web server must be established. At present, no such server exists for this project. We hope to be able to implement one in the future. Such a server could serve two main purposes. First, it can provide a tool for a biologist to input a sequence of interest and obtain a confidence score that might aid them in prioritizing wet-lab experiments. Second, it could display the evidence extracted from biomedical literature about the input sequence, so that the biologist has the opportunity to validate the predictions which could then be fed as input to a crowd-sourced biocuration effort on ACPs.

PREDICTING PROTEIN FUNCTION

Introduction

Proteins can have many properties. A system known as the Gene Ontology (GO) seeks to classify these properties hierarchically [1]. For instance, the trait “DNA-binding transcription factor activity” is a child of “Transcription regulator activity,” which is itself a child of “Molecular function.” The ontology has three roots: “Biological process,” “Cellular component,” and “Molecular function.” The Critical Assessment of protein Function Annotation (CAFA) [31] competition invites researchers to computationally identify GO annotations for a given protein. Once the competition is over, researchers perform lab work to confirm the annotations. The CAFA administrators then compare the predictions from the competition to the results from these experiments to identify whose models performed the best [31].

As described in the previous chapter, *DeepACPpred* is a system which identifies whether or not a sequence of amino acids has anticancer properties. At a high level, the specific property being trained for is not important; this system could be used to train for any property or set of properties. Thus, we adapted *DeepACPpred* model for predicting GO annotations to generate submissions for this year’s CAFA4¹.

While the official CAFA4 results are still pending, we internally evaluated our models using the publicly available CAFA3 benchmark (test) data. While the results produced were promising, over the course of this project, we discovered that there are some issues in applying *DeepACPpred* model to this problem. Specifically, because proteins are longer than peptides, the amount of time and resources required for training are increased dramatically. Nonetheless, this system does show promise, and

¹<https://www.biofunctionprediction.org/cafa/>

these issues are not prohibitive to its refinement for future CAFA submissions.

Related Work

A number of other projects have used deep learning to predict protein function [2, 15, 16]. One of these systems, *ProLanGO* by Cao et al. [2], also used RNNs to do this. Their method identified all k -mers (subsequences of length k) of amino acids that had a frequency above some threshold, and where k was 3, 4, or 5. These subsequences were considered “words,” common units within proteins. Each sequence was then split into the words, with preference given to longer ones. Unknown sequences were kept as an unidentified word. They called this input sequence “ProLan.” They then converted each GO term into a sequence of 4 letters (with 26^4 possible sequences, this exceeds the number of terms in GO), and called this set of sequences the “GOLan” language. Thus, they are able to treat the problem as language translation.

To perform this task, they use neural networks. The first layer is an RNN, which converts each word into a vector. The second RNN then performs the actual task of translation. Their training set consisted of 419,192 amino acid sequences from all species, and their testing set had 104,798 amino acid sequences. With such a large dataset, they were able to train this neural network over 78,000 steps. Ultimately, their results were comparable to the state-of-the-art methods [2].

Another project, *DeepGO* by Kulmanov et al. [16] used convolutional neural networks to perform this task. Here, the first layer was an embeddings layer; however, unlike *DeepACPPred*, the inputs to this layer were the trigrams of the amino acid sequence. This means that there were 8000 (20^3) elements in the vocabulary for the embeddings layer. After this was a convolutional neural network with max pooling applied. This is also concatenated with PPI (protein-protein interaction) data.

The final step is a series of fully-connected layers that form a directed acyclic

graph (DAG) following the structure of the Gene Ontology. Thus, there is one fully-connected layer per GO class. The output of these layers corresponds to a confidence score for the associated GO annotation. *DeepGO*'s performance was also comparable to that of other state-of-the-art methods [16].

The authors of *DeepGO* later improved its performance with *DeepGOPlus* [15]. This project featured three key differences from *DeepGO*. First, rather than use an embeddings layer, they used a simple one-hot vector encoding for their model. This was chosen because they found that the embeddings would often greatly contribute to overfitting. Second, they added multiple CNNs, each with a different number of filters. Finally, rather than use the hierarchical classification scheme outlined above, they used a simple flat classification layer like the one presented here. This was chosen to limit memory usage [15].

Methods

Datasets

The first step was to gather datasets that are applicable to this specific problem. For CAFA4, we were asked to predict GO annotations for a number of proteins from a number of species. It was decided that a training set would be created for each species, as it was posited that proteins may have different roles in different species. Further, by differentiating by species, we were able to perform optimization within a reasonable amount of time. To do this, every protein with GO annotations verified by wet-lab experiments for each species was downloaded from UniProtKB². The target proteins specified by CAFA organizers were filtered out.

Further, if an annotation (i.e. GO category) appeared in fewer than 1% of all

²<https://www.uniprot.org/help/uniprotkb>

proteins, then it was removed from the possible outputs. This was done because annotations that do not appear frequently are difficult to predict. If the system were to predict that all given proteins did not have the given annotation, there would be a large number of true negatives and a small number of false negatives. The loss produced by those false negatives is relatively negligible; there is little incentive for the gradient descent algorithm to optimize for this annotation in comparison to the more frequent annotations.


Adapting *DeepACPpred*

Because of the similarities between the problems of ACP prediction and GO annotation prediction, very little had to be changed to adapt *DeepACPpred* model for this task. The first major difference is the removal of transfer learning. Because the datasets were all quite large, training on a larger dataset before the given training set was unnecessary.

The second major difference comes in the output of the model. In *DeepACPpred*, the expected result is simply a single value indicating its confidence that the given amino acid sequence was an ACP. For GO annotation prediction, there is one confidence score for each annotation being considered. Thus, the number of nodes in the final layer is equal to the number of annotations. Additionally, it was discovered that the presence of the dense layer between the concatenation and output layers (See figure 2.3) decreased performance, so it was removed.

Training

To find the optimal parameters for the neural network, ten fold cross-validation was used on the datasets for humans and *Pseudomonas aeruginosa*, a species of bacteria. This was done because it was theorized that since prokaryotes are generally much simpler than eukaryotes, different networks should be used for each kind of



(a) Lengths of proteins for human


(b) Lengths of proteins for *E. coli*

Figure 3.1: A histogram comparison of the sequence lengths in a selection of unfiltered datasets.

organism. During this process, several discoveries were made.

First, because proteins are longer sequences of amino acids than peptides, a larger number of parameters was required to prevent a vanishing gradient. This may be because the LSTMs needed more variables to “remember” data learned in previous states. Put another way, with a higher amount of data, more information is retained in later states. Figure C.1 shows the lengths of the sequences from humans and *E. coli*. For data from all relevant species, see appendix C. These distributions are skewed extremely right.

Regardless of the reason for the increased parameter requirements, the increased network size meant that the amount of time required for training increased to the point of being infeasible. For instance, on unfiltered human proteins, training for a single epoch could take multiple hours depending on the exact parameters of the network. This makes large-scale optimization extremely difficult. Worse, the amount of memory required for these networks exceeded what was on the available GPU (12GB). Thus, measures had to be taken to reduce the number of variables in these networks. First, proteins longer than 900 amino acids were filtered out. With the shorter proteins, fewer variables were needed in each LSTM. Since the distribution



(a) Filtered lengths of proteins for human

(b) Filtered lengths of proteins for *E. coli*

Figure 3.2: A histogram comparison of the sequence lengths in a selection of filtered datasets.

of lengths of proteins was skewed so dramatically to the right, relatively few proteins were removed from this cutoff. Figure C.2 shows the lengths of the proteins in the resulting datasets for humans and *E. coli*. See appendix C for data on the datasets for all relevant species.

Additionally, for the training sets with many proteins, a random sample of 2,000 was taken for training. While not ideal, this was a necessary step to facilitate the training process. This random sample was not used when using the CAFA 3 data as a testing set. Finally, we limited the datasets to only train and test on terms within the “Molecular Function” subontology. This was chosen because the RNN structure of the neural network was designed around the molecular structure of proteins.

Training was only performed on leaf terms, i.e., ones that had no children. Before evaluation, the predictions were propagated through the GO annotation DAG. The result is that each node’s predicted confidence score was equal to the greatest of its children. This is the same process used by the organizers of CAFA, and it has the ultimate effect of boosting performance, since the upper levels of the ontology have fewer terms. This means that if just one of an internal node’s children is predicted to be true, then it will itself also be predicted to be true. Since the false negatives greatly

outweigh the false positives, having more positive predictions benefits the network.

Results and Discussion

The official results of CAFA 4 will not be made available until later this year. While we cannot measure performance for the CAFA 4 test data, we can measure against the data from CAFA 3. To produce test results, the network was trained on a dataset with the targets from CAFA 3 filtered out. Then, the targets were fed into the network, and the results were compared against those of CAFA 3. Some species present in CAFA 4 are not in CAFA 3, and vice versa. Only the species from CAFA 3 are shown here. Additionally, if a species had no annotations for any of its proteins after removing the ones that appear for fewer than 1% of proteins in the training set, it was skipped entirely.

CAFA 3 evaluated the results in a number of different ways. First, they made a distinction between proteins that had had no GO annotations prior to the submission deadline and proteins that had had some GO annotations. We used the “No knowledge” model, as that was also the one that the CAFA organizers reported in the primary portion of their paper; the results from the “Limited” data were reported with the supplementary material [31]. Additionally, some models would not make any predictions for a protein. In the “partial” evaluation mode, these are removed from consideration. In the “full” mode, they are considered, and likely reduced the values of the performance metrics. We used the “full” mode, since, again, the CAFA organizers reported this in the main portion of the paper [31].

We had two separate neural network models at the end of this process: a “large” one from optimization on the human data, and a “small” one from optimization on the bacteria. Table 3.1 shows the F_{max} score produced by each network for each species. The F_{max} is found by calculating the F1 score at numerous thresholds (a threshold

being the confidence score above which a prediction is classified as a positive and below which as a negative) and taking the maximum of these. Additionally, the results of 5-fold cross validation for each network type and for each species is shown here.

While on the surface these F_{max} values are much lower, especially when compared to those of ACP prediction, it is important to remember that the F_{max} score of the best-performing model of CAFA 3 for the molecular function subontology was about 0.6 [31]. Even with this in mind, however, the results are certainly suboptimal.

While analyzing this data, we found that the “small” neural network seemed to outperform the “large” one for nearly every species. This contradicted our initial hypothesis, that the “large” model would work better for the more complicated eukaryotes and the “small” model would work better for the simpler prokaryotes. This hypothesis was based on the idea that in more complicated species, proteins might have more varied roles and interact with each other in ways not seen in simpler organisms. These results seem to indicate that the complexity of the organism does not affect the size of network needed to predict its GO terms accurately.

It is also worth noting that the cross validation performed much better than the testing procedure. A previous study by Kahanda et al. demonstrated this as well [14]. The worse performance may be due to the fact that the GO terms on the proteins in the testing set are the results of new experiments. It is likely that the first sets of GO annotations for any species are the most common ones, which should be easier to predict. Under this assumption, the new GO terms would be rarer and thus more difficult to predict. Kahanda et al. reached a similar conclusion. [14].

When comparing to the other deep learning methods discussed above, only *DeepGO* used the exact same evaluation procedures. Their F_{max} score of 0.47 across all species when evaluating on the CAFA 3 dataset beats our model’s performance [16],

Species	# Train	# Test	L, Test	S, Test	L, CV	S, CV
Mouse-ear Cress	9209	88	0.269	0.283	0.624	0.614
Fruit Fly	1867	46	0.202	0.230	0.487	0.498
Zebrafish	1366	4	0.300	0.316	0.501	0.501
African clawed frog	1710	5	0.143	0.229	0.588	0.610
Human	11040	752	0.245	0.259	0.478	0.461
Mouse	9465	289	0.260	0.248	0.422	0.420
Rat	4835	48	0.378	0.391	0.407	0.412
Slime mold	1654	13	0.279	0.529	0.305	0.531
<i>E. coli</i>	2237	106	0.249	0.271	0.503	0.517
<i>H. pylori</i>	353	1	0.160	0.160	0.591	0.585
<i>P. putida</i>	414	1	0.194	0.194	0.587	0.604
<i>B. subtilis</i>	1864	7	0.185	0.202	0.502	0.512
<i>S. solfataricus</i>	336	2	0.207	0.323	0.519	0.536
<i>S. pombe</i>	2250	59	0.201	0.202	0.442	0.438
<i>S. cerevisiae</i>	2705	69	0.316	0.299	0.432	0.437
Average	3,420	99	0.221	0.262	0.493	0.512

Table 3.1: F_{max} scores of this model when testing on CAFA 3 result data and performing 5-fold cross validation with training data.

but is also beat by the best-performing model presented in CAFA 3 (whose F_{max} score, as stated above, was 0.6) [31].

The researchers who created *ProLanGO*, meanwhile, reported precision and recall, metrics used as intermediates for calculating the F1 score, at various thresholds; however, they only report this for the validation sets and not for testing on the CAFA 3 results. From their data, we can see that they had an F_{max} score of about 0.725 on the validation sets. When comparing our validation results to theirs, it is clear that their model is superior. It is worth noting that these researchers reported superior performance from *DeepGO* vs *ProLanGO*.

Conclusion and Future Work

The results shown here for using *DeepACPred* to predict protein annotations are somewhat disappointing. We have a couple of theories as to why this may be the case. To begin with, the longer sequence length of proteins necessitates the use of larger LSTM structures or advanced techniques not explored in this project, such as bucketing. Our hardware was unable to handle the size of the neural network truly necessary to achieve excellent results. This is likely primarily due to the convolution layer, which can be about ten times as large as any of the other layers. Comparatively, *ProLanGO* did not use a convolutional layer, and so they likely did not face the memory problems that we had [2].

Additionally, one of our assumptions was that each species should be trained separately. This was based on the idea that proteins may have different roles in each organism or interact with its other proteins in unique ways. However, these effects may have been minimal overall. It may have been more beneficial to train on the larger dataset of all experimentally-verified proteins to help prevent overfitting, as was done in *ProLanGO* [2]. This would have been especially helpful for a larger

neural network. It is important to remember, however, that one of the justifications for using this technique was the long training times. By separating the species, we were able to perform optimization experiments within a reasonable amount of time. Regardless, without accounting for time and hardware restrictions, the best technique may have been to use the full dataset.

We conclude that while *DeepACPred* is a model that works well with peptides, it has structural problems that currently prevent its use with full proteins. In particular, the optimal convolutional network may simply be too large for the resources available to us. With better hardware, it may prove to give better performance than what is seen here.

CONCLUSION AND FUTURE WORK

Over the course of this project, we created a hybrid CNN-RNN machine learning model to predict the properties of a given amino acid sequence. Chapter 2 analyzed this system’s applicability to the problem of predicting anticancer peptides (ACPs) and found results comparable to those of the state-of-the-art methods. As *DeepACPred* is among the first RNN-based models for this problem, our hope is that future work could potentially improve upon our results. In particular, we hope that as available ACP data grows larger, the performance of deep learning systems will outperform that of other machine learning methods. Further, we identified potential additional parameters that could be used to bolster performance.

In chapter 3, we adapted the *DeepACPred* model to the problem of identifying a protein’s Gene Ontology (GO) annotations for submission to the CAFA competition. We discovered that the system had worse performance here. We identified many potential reasons for this, but the primary one is the fact the longer amino acid sequences means that training will require greater resources, both in terms of hardware and time. The biggest contributor to this factor is the CNN layer, which can have as many as 10 times as many parameters as the other layers. The end result is that we were forced to use smaller, suboptimal networks. Ultimately, with better hardware and more time, this model could potentially be useful for this problem.

In summary, we set out to show that an RNN-based system could be used to predict the properties of amino acid sequences. Specifically, our initial focus was on the prediction of ACPs. The system we developed, *DeepACPred* was a novel hybrid CNN-RNN model. Unfortunately, while trying to adapt the underlying model for the task of predicting protein function, we noticed that it struggled with longer proteins, likely because it was developed specifically to predict the properties

of shorter peptides. *DeepACPpred* ultimately got excellent results for ACPs, and we hope that future development may build upon what we have shown here.

REFERENCES CITED

- [1] Michael Ashburner, Catherine A. Ball, Judith A. Blake, David Botstein, et al. Gene Ontology: tool for the unification of biology. *Nature Genetics*, 25(1):25–29, 2000.
- [2] Renzhi Cao, Colton Freitas, Leong Chan, Miao Sun, et al. ProLanGO: Protein function prediction using neural machine translation based on a recurrent neural network. *Molecules*, 22(10), 2017.
- [3] Wei Chen, Hui Ding, Pengmian Feng, Hao Lin, and Kuo-Chen Chou. iACP: a sequence-based tool for identifying anticancer peptides. *Oncotarget*, 7(13):16895–16909, Mar 2016. 26942877[pmid].
- [4] Alan Coates, Suzanne Abraham, S.B. Kaye, Timothy Sowerbutts, et al. On the receiving end—patient perception of the side-effects of cancer chemotherapy. *European Journal of Cancer and Clinical Oncology*, 19(2):203 – 208, 1983.
- [5] Teerasak E-Kobon, Pennapa Thongararm, Sittiruk Roytrakul, Ladda Meesuk, and Pramote Chumnannpuen. Prediction of anticancer peptides against MCF-7 breast cancer cells from the peptidomes of *Achatina fulica* mucus fractions. *Computational and structural biotechnology journal*, 14:49–57, Dec 2015. 26862373[pmid].
- [6] Diana Gaspar, A. Salomé Veiga, and Miguel A.R.B. Castanho. From antimicrobial to anticancer peptides. A review. *Frontiers in Microbiology*, 4:294, 2013.
- [7] Francesca Grisoni, Claudia S. Neuhaus, Gisela Gabernet, Alex T. Müller, et al. Designing anticancer peptides by constructive machine learning. *ChemMedChem*, 13(13):1300–1302, 2018.
- [8] Francesca Grisoni, Claudia S. Neuhaus, Miyabi Hishinuma, Gisela Gabernet, et al. De novo design of anticancer peptides by ensemble artificial neural networks. *Journal of Molecular Modeling*, 25(5):112, 2019.
- [9] Zohre Hajisharifi, Moien Piryaiee, Majid Mohammad Beigi, Mandana Behbahani, and Hassan Mohabatkar. Predicting anticancer peptides with Chou’s pseudo amino acid composition and investigating their mutagenicity via Ames test. *Journal of Theoretical Biology*, 341:34 – 40, 2014.
- [10] Md-Nafiz Hamid and Iddo Friedberg. Identifying antimicrobial peptides using word embedding with deep recurrent neural networks. *Bioinformatics*, 35(12):2009–2016, 11 2018.
- [11] Frederick Harris, Sarah R. Dennison, Jaipaul Singh, and David A. Phoenix. On the selectivity and efficacy of defense peptides with respect to cancer cells. *Medicinal Research Reviews*, 33(1):190–234, 2013.

- [12] David W. Hoskin and Ayyalusamy Ramamoorthy. Studies on anticancer activities of antimicrobial peptides. *Biochimica et Biophysica Acta (BBA) - Biomembranes*, 1778(2):357 – 375, 2008.
- [13] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift, 2015.
- [14] Indika Kahanda, Christopher S. Funk, Fahad Ullah, Karin M. Verspoor, and Asa Ben-Hur. A close look at protein function prediction evaluation protocols. *GigaScience*, 4(1):41, 2015.
- [15] Maxat Kulmanov and Robert Hoehndorf. DeepGOPlus: improved protein function prediction from sequence. *Bioinformatics*, 36(2):422–429, 07 2019.
- [16] Maxat Kulmanov, Mohammed Asif Khan, and Robert Hoehndorf. DeepGO: predicting protein functions from sequence and interactions using a deep ontology-aware classifier. *Bioinformatics*, 34(4):660–668, 10 2017.
- [17] Sneha Lata, Nitish K. Mishra, and Gajendra P. S. Raghava. AntiBP2: improved version of antibacterial peptide prediction. *BMC bioinformatics*, 11 Suppl 1(Suppl 1):S19–S19, Jan 2010. 20122190[pmid].
- [18] DB Longley and PG Johnston. Molecular mechanisms of drug resistance. *The Journal of Pathology*, 205(2):275–292, 2005.
- [19] Balachandran Manavalan, Shaherin Basith, Tae Hwan Shin, Sun Choi, et al. MLACP: machine-learning-based prediction of anticancer peptides. *Oncotarget*, 8(44):77121–77136, 2017.
- [20] Prabina Kumar Meher, Tanmaya Kumar Sahu, Varsha Saini, and Atmakuri Ramakrishna Rao. Predicting antimicrobial peptides with improved accuracy by incorporating the compositional, physico-chemical and structural features into Chou’s general PseAAC. *Scientific Reports*, 7(42362), 2017.
- [21] Nishant Thakur, Abid Qureshi, and Manoj Kumar. AVPpred: collection and prediction of highly effective antiviral peptides. *Nucleic acids research*, 40(Web Server issue):W199–W204, Jul 2012. 22638580[pmid].
- [22] Atul Tyagi, Pallavi Kapoor, Rahul Kumar, Kumardeep Chaudhary, et al. In Silico Models for Designing and Discovering Novel Anticancer Peptides. *Scientific Reports*, 3(2984):2045–2322, 2013.
- [23] Atul Tyagi, Abhishek Tuknait, Priya Anand, Sudheer Gupta, et al. CancerPPD: a database of anticancer peptides and proteins. *Nucleic acids research*, 43(Database issue):D837–D843, Jan 2015. 25270878[pmid].

- [24] Guangshun Wang, Xia Li, and Zhe Wang. APD2: the updated antimicrobial peptide database and its application in peptide design. *Nucleic Acids Research*, 37(suppl_1):D933–D937, 10 2008.
- [25] Guangshun Wang, Xia Li, and Zhe Wang. APD3: the antimicrobial peptide database as a tool for research and education. *Nucleic acids research*, 44(D1):D1087–D1093, Jan 2016. 26602694[pmid].
- [26] Leyi Wei, Yijie Ding, Ran Su, Jijun Tang, and Quan Zou. Prediction of human protein subcellular localization using deep learning. *Journal of Parallel and Distributed Computing*, 117:212 – 217, 2018.
- [27] Chuanyan Wu, Rui Gao, Yusen Zhang, and Yang De Marinis. PTPD: predicting therapeutic peptides by deep learning and word2vec. *BMC Bioinformatics*, 20(1):456, 2019.
- [28] Hai-Cheng Yi, Zhu-Hong You, Xi Zhou, Li Cheng, et al. ACP-DL: A Deep Learning Long Short-Term Memory Model to Predict Anticancer Peptides Using High-Efficiency Feature Representation. *Molecular therapy. Nucleic acids*, 17:1–9, Sep 2019. 31173946[pmid].
- [29] Z. You, M. Zhou, X. Luo, and S. Li. Highly efficient framework for predicting interactions between proteins. *IEEE Transactions on Cybernetics*, 47(3):731–743, March 2017.
- [30] Xiaowei Zhao, Hongyu Wu, Hairong Lu, Guodong Li, and Qingshan Huang. LAMP: A database linking antimicrobial peptides. *PloS one*, 8(6):e66557–e66557, Jun 2013. 23825543[pmid].
- [31] Naihui Zhou, Yuxiang Jiang, Timothy R. Bergquist, Alexandra J. Lee, et al. The CAFA challenge reports improved protein function prediction and new functional annotations for hundreds of genes through experimental screens. *Genome Biology*, 20(1):244, 2019.

APPENDICES

APPENDIX A

PROTEINS

As this paper deals largely with the application of computing techniques to the problem of identifying qualities of proteins, it is important to understand the nature of these chemicals. This section discusses the structure and biological importance of proteins.

In summary, proteins are polymers, chemicals made of chains of smaller units called monomers. For proteins, these monomers are a class of chemicals known as the amino acids. When many amino acids are chained together in a sequence, weak chemical interactions cause them to fold into more complicated structures. Thus, they are able to perform a great many tasks, from holding oxygen in red blood cells to facilitating muscle movement.

Amino Acids

Amino acids are the chemicals which form the basis for proteins. Each one contains a carboxyl group (-COOH) and an amine group (-NH₂). Etymologically, “amine” has its roots in the word “ammonia,” a chemical with a very similar structure (NH₃). An amine group is just an ammonia molecule with a hydrogen removed. Between the carboxyl and amine groups is an R group, otherwise known as a side chain. The exact chemical structure of the R group varies among amino acids. Figure A.1 shows this chemical structure.




Figure A.1: Basic amino acid structure.

When two amino acids combine in the formation of a peptide bond, the hydroxide group (-OH) from one of the amino acid’s carboxyl group detaches and forms a water molecule (H₂O) with one of the hydrogens (H) from the other’s amine group. This leaves a site on each amino acid where a bond can be formed. This whole reaction, since it generates a water molecule, is called a condensation reaction, and it is shown in Figure A.2.

While there are 22 amino acids, two of them, selenocystine and pyrrolysine, are rare and do not appear in any of the peptides or proteins analyzed in this project. The rest, along with their 3-letter and 1-letter abbreviations, are shown in table A.1.




Figure A.2: Condensation reaction.

Protein Folding

The sequence of amino acids that creates a protein is known as its primary structure. As each protein is constructed, weak bonds and chemical interactions among the amino acids' R groups forces them to fold into complex shapes. The simplest of these are β -sheets, flat surfaces, and α -helices, twisting structures. These are collectively known as the secondary structure of a protein. A protein's secondary structure is shown in figure A.3.

These come together to for the protein's final, tertiary, structure, shown in figure A.4. At this point, with no external influence, there is no higher structure. It is possible, however, for peptide chains to interact and form even larger structures. These are known as a quaternary structure, shown in figure A.5.

Ultimately, this project only analyzed the primary structure of each peptide or protein. This is a key determining factor in its structure, and therefore function, but certainly not the only one, as the existence of quaternary structure shows.

¹<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

Name	3-letter	1-letter
Alanine	Ala	A
Arginine	Arg	R
Asparagine	Asn	N
Aspartic acid	Asp	D
Cysteine	Cys	C
Glutamine	Gln	Q
Glutamic acid	Glu	E
Glycine	Gly	G
Histidine	His	H
Isoleucine	Ile	I
Leucine	Leu	L
Lysine	Lys	K
Methionine	Met	M
Phenylalanine	Phe	F
Proline	Pro	P
Serine	Ser	S
Threonine	Thr	T
Tryptophan	Trp	W
Tyrosine	Tyr	Y
Valine	Val	V

Table A.1: Amino acids with their abbreviations




Figure A.3: Secondary structure of a protein. Image by Thomas Shafee. Licensed under the Creative Commons Attribution-Share Alike 4.0 International license¹




Figure A.4: Tertiary structure of a protein. Original image by Thomas Shafee and cropped to show only the relevant elements. Licensed under the Creative Commons Attribution-Share Alike 4.0 International license

Quaternary




Figure A.5: Quaternary structure of a protein. Image by Thomas Shafee and cropped to show only the relevant elements. Licensed under the Creative Commons Attribution-Share Alike 4.0 International license

APPENDIX B

NEURAL NETWORKS

A neural network contains many nodes, called neurons, which contain values. These nodes are organized into layers, and connections are made between the layers. There are many kinds of layers, each with different behaviors. The ones used in this project are outlined in the sections of this appendix.

Dense Layers

The simplest layer is known in *Tensorflow* as a Dense layer. Here, each node is connected to every node in the layer before it. Figure B.1 shows a small example of such a layer.




Figure B.1: A simple example of a neural network

Each connection multiplies the value in its root node by some weight. The destination node sums the values coming into it from each edge. In a mathematical equation, if x is an indexable vector of size n representing the values in the input layer, y is a similarly indexable vector of size m representing the values in the output layer, and W is a two-dimensional matrix representing the weights of each edge, then

$$y_i = \sum_{j=0}^m W_{ij} * x_j \quad (\text{B.1})$$

for $0 \leq i \leq n$. Ultimately, this is identical to multiplying the vector x by the matrix W . An additional bias value may be added to the output, giving a final equation of

$$y = Wx + b \quad (\text{B.2})$$

Thus, the most basic neural network layer can be represented by a simple mathematical equation.

Activation

Many times, it is desired that the range of the outputs of a layer be restricted to some range. In this case, an activation function is applied to each output. There are many kinds of activation functions, but only the sigmoid function was used in this project.

The output of the final layer of *DeepACPpred* is desired to be within the range $[0, 1]$. Any intermediate value is rounded to 0, which indicates that the input amino acid sequence was not an ACP, or 1, which indicates that it was an ACP. The sigmoid function has the range $(0, 1)$, and is therefore appropriate for this purpose. It is given by the following equation:

$$\sigma(t) = \frac{e^t}{e^t + 1}. \quad (\text{B.3})$$

Its plot is shown in figure B.2.




Figure B.2: Plot of the sigmoid function. Image by Wikimedia user Qef and released to the public domain.

Loss

In order for a neural network to be trainable, a loss function must be applied to its output. This function indicates how big of a difference there is between the expected output of the network and its true output. The loss function used in this project was the binary cross entropy function.

The i th input has an expected output value that can be given by y_i , which will be either a 0 or a 1. The confidence score produced by the network is given by p_i . The binary cross entropy produced for the i th value is given as

$$Loss = \begin{cases} -\log(p_i), & \text{if } y_i = 1 \\ -\log(1 - p_i), & \text{if } y_i = 0 \end{cases} \quad (\text{B.4})$$

$$= -y_i * \log(p_i) - (1 - y_i) * \log(1 - p_i) \quad (\text{B.5})$$

The final value used by the network for optimization is the average of all of these losses.

Gradient Descent

The weights and biases of a network are collectively known as its parameters. When we calculate the loss of a network given a set of inputs, we can consider the parameters to be variables. To improve the performance of the network (i.e., minimize loss), we must adjust these parameters. This is done via gradient descent.

If we consider the parameters to be the variables of the loss function, we can differentiate with respect to any of these. If we do this for every parameter, the result is a vector known as the gradient.

$$\nabla l(w_0, \dots, w_n, b_0, \dots, b_m) = \left\langle \frac{\partial l}{\partial w_0}, \dots, \frac{\partial l}{\partial w_n}, \frac{\partial l}{\partial b_0}, \dots, \frac{\partial l}{\partial b_m} \right\rangle \quad (\text{B.6})$$

When the current parameters are applied to the gradient, the result is the direction of greatest ascent. The negative is the direction of greatest descent. Multiplying the gradient by some constant called the learning rate η (which could be considered the size of the “step” we are taking) and then subtracting it from the vector of parameters results in an adjusted parameter field which should have a smaller loss than before. Repeating this process minimizes the loss function.

$$\langle w_0, \dots, w_n, b_0, \dots, b_m \rangle = \langle w_0, \dots, w_n, b_0, \dots, b_m \rangle - \eta \nabla l(w_0, \dots, w_n, b_0, \dots, b_m) \quad (\text{B.7})$$

Embeddings

A common way to represent each element in a sequence of a neural network is with embedding vectors. Ideally, these vectors will be similar for similar inputs. As an example, for amino acids, asparagine and glutamine have similar chemical properties, so their embeddings should also be similar. As another example in the realm of language, the embedding for the word “king” minus that for “man” should be a vector containing the concept of “royalty.” When added to “woman,” the resulting vector should be close to “queen.”

$$E(\textit{“queen”}) \approx E(\textit{“king”}) - E(\textit{“man”}) + E(\textit{“woman”}) \quad (\text{B.8})$$

These embeddings can be generated by the neural network during the learning process. As this is the first layer of the neural network in this project (and indeed, in many networks in which it appears), the input for the network as a whole must be customized for this layer in particular. Each element of the sequence must be converted to a corresponding integer before training. When the layer is initialized, a matrix whose width is the size of the vocabulary and height is the specified size for the embeddings is created. When the data is input into the layer, each integer is converted into a one-hot vector, i.e., a vector whose length is the size of the vocabulary and is initialized to all 0s except for a single 1. If the integer in question is i , then the i th element in the one-hot vector is the 1. The one-hot vector is then multiplied by the matrix described above. This has the ultimate effect of selecting the i th column of the matrix. If n is the size of the vocabulary, m is the length of embedding vector that is desired, and we are looking for the second embedding, then

$$E(2) = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,n} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m,1} & a_{m,2} & \cdots & a_{m,n} \end{pmatrix} * \begin{pmatrix} 0 \\ 1 \\ \vdots \\ 0 \end{pmatrix} = \begin{pmatrix} a_{1,2} \\ a_{2,2} \\ \vdots \\ a_{m,2} \end{pmatrix} \quad (\text{B.9})$$

Throughout the learning process, the values in the matrix described above are treated as parameters, and are therefore adjusted. Since an embeddings layer converts a single value to a one-dimensional vector, its output has one additional dimension to the input.

Word2vec

It is possible to use embedding vectors that were created prior to training. One algorithm to pre-train embeddings is Word2vec. There are two methods within this algorithm: skip-gram and continuous bag-of-words (CBOW).

In skip-gram, a neural network is created which has as input a one-hot vector representation of the word or token whose vector we want to generate. The output

is a bag-of-words vector representation of the tokens within a certain distance of the target. This bag-of-words vector has a length equal to the number of tokens in the vocabulary. Each element in the vector represents a count of how many times a token appears within the current frame (oftentimes with more distant tokens being weighted less than near ones). Between the input and output is a Dense neural network layer whose size is equal to the desired length of the embedding vector. This neural network is trained for every token in a given sequence. The embedding for a token is then the vector contained in the hidden layer when that token's one-hot vector is used as input to the neural network.

$$\begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \rightarrow \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} \rightarrow \begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_V \end{pmatrix}$$

Figure B.3: Representation of skip-gram

Thus, it could be said that the skip-gram model matches a token to its context. The inverse is true for the CBOW model. This model uses as input the bag-of-words for the tokens surrounding a target, and as output the target word itself.

$$\begin{pmatrix} c_1 \\ c_2 \\ \vdots \\ c_V \end{pmatrix} \rightarrow \begin{pmatrix} a_1 \\ a_2 \\ \vdots \\ a_n \end{pmatrix} \rightarrow \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

Figure B.4: Representation of CBOW

Recurrent Neural Networks

A recurrent neural network (RNN) applies an operation repeatedly to a series of inputs. It does this by combining the output of the previous iteration with the next

input vector and feeding this back into itself, forming a form of feedback. The input to an RNN must be two-dimensional, such as a sequence of n vectors of size m .

If the output size of the RNN is h , then on the first iteration, the RNN will take the first element in the input sequence and concatenate it with h 0s (Since no output has been produced yet, the network must be seeded with these 0s). The layer will then output a vector of size h . The network will then concatenate this with the second element in the input sequence and apply the operation again. This repeats until the network reaches the final element of the input, at which point, the vector of size h is provided as output for the rest of the network. This process is shown in figure B.5.




Figure B.5: Diagram showing how an RNN works. Image created by Wikimedia user fdeloche and licensed under the Creative Commons Attribution-Share Alike 4.0 International license.

The nature of the operation that is performed depends on what kind of RNN is being used. For a simple RNN, this is identical to the operation of a Dense neural network that takes an input of size $h + m$ and has an output of size h . More complicated RNNs are discussed in future sections.

Long Short-Term Memory Layers

Long Short-Term Memory layers (LSTMs) are a more advanced form of an RNN. Standard RNNs often face exploding or vanishing gradients (see sections below). To remedy this, LSTMs provide a much more complicated structure with an additional vector, called the cell state, of size $h + m$ and a number of “gates” that determine what information is kept, removed, or added to that cell state. Figure B.6 shows this structure.

The first gate is called the “forget gate” and controls what information in the cell state is dropped. It concatenates the output of the previous cell with the next input vector, then applies a dense neural network with a sigmoid activation function




Figure B.6: Diagram showing how an LSTM works. Orange sections are dense networks, and yellow sections signify piecewise operations. Image created by Wikimedia user Guillaume Chevalier and licensed under the Creative Commons Attribution-Share Alike 4.0 International license. Edited to show relevant portions.

to the result. This provides a vector of size $h + m$ whose values all lie in the range $(0, 1)$. This is then multiplied piecewise with the cell state. If a value in the output of the dense network is close to 0, then the effect is that the associated value in the cell state becomes 0 and is “forgotten.” Likewise, if the value was close to 1, then the value in the cell state remains the same and is kept.

The second gate is the “input gate” and controls what new information is placed in the cell state. First, two dense networks are applied to the concatenated input. Once one of these is activated with the sigmoid function, and the other is activated with the $\tanh$ function, which has a range of $(-1, 1)$. The results are then multiplied piecewise, and as before, this has the effect of causing values in the $\tanh$ vector to be dropped or kept. Put another way, the values in the $\tanh$ vector are “candidates” to be added to the cell state, and the values in the sigmoid vector determine which ones are kept and which ones are excluded. The final result of this gate is added piecewise to the cell state.

The final gate is the “output gate” and controls what elements in the cell state are provided as the output of the cell. First, another dense layer with a sigmoid activation function is applied to the concatenated input. This is multiplied by the cell state vector, which has had a $\tanh$ activation function applied to it. As before, this has the effect of selecting which values in the activated cell state are provided as the final output of the cell.

Convolutional Networks

A convolution applies some kind of filter to a sliding window of the elements of a list. In the context of a neural network, this filter has a vector whose length is the size of the window. It then takes the values in the window and puts them in a vector which is dotted with the filter vector. An example convolution is shown below.

$$\begin{aligned}\langle 5, 1, 3, 7, 0 \rangle * \langle 2, 0, 1 \rangle &= \langle \langle 5, 1, 3 \rangle \cdot \langle 2, 0, 1 \rangle, \\ &\quad \langle 1, 3, 7 \rangle \cdot \langle 2, 0, 1 \rangle, \\ &\quad \langle 3, 7, 0 \rangle \cdot \langle 2, 0, 1 \rangle \rangle \\ &= \langle 13, 9, 6 \rangle\end{aligned}$$

This kind of layer is commonly used in image processing. It is particularly good at identifying low-level details in images and assembling them together. In this project, we’re hoping the the layer can identify chemical reactions among nearby amino acids.

Vanishing Gradients and Overfitting

When training a neural network, some problems may arise which hinder its learning capabilities. The first of these that we will discuss is vanishing gradients. During gradient descent, it is possible that a sub-optimal local minimum may be found. In this case, the gradient will be the 0 vector, and so no improvements to loss will be made. During this project, this was seen especially when training on the proteins. We theorize that since LSTMs need larger cell states to operate efficiently on longer sequences, training on proteins will require many more parameters than peptides for proper training.

The second major issue is that of overfitting. This is when a neural network has significantly better performance on a training set than on a validation or testing set. This happens because the network is effectively “memorizing” the elements in the training set, and no longer learning the fundamental qualities of the data in the input.

To remedy this, several techniques can be applied. First, overfitting often occurs when there are too many parameters in a neural network. These cannot always be reduced, however, as having too few parameters can lead to vanishing gradients. A second way to reduce overfitting is to increase the size of the training set. It is harder for the network to “memorize” a larger input set. This is also not always possible, as is the case in the ACP prediction problem, as there are not many experimentally-verified ACPs. Two additional techniques are dropout, regularization, and batch normalization, each of which have their own sections.

Dropout

Dropout is a technique to reduce the effects of overfitting. This is done by randomly disconnecting a given proportion of the connections between layers. The idea is that hopefully the elements within the network should be able to compensate for the absence of data. Metaphorically, if some proportion of a company's workers were randomly selected to be absent each day, the remaining workers would have to know what everyone else was doing.

Application of this technique can often require a slight increase in the size of the network.

Regularization

Sometimes overfitting happens when the values within a network become too extreme and introduce a high level of complexity. Figure B.7 shows the difference between a simple function and a complex one. A function with higher complexity has more variance and is therefore less likely to be suitable for data outside the given training set.

Regularization puts a restriction on the range of the parameters of a network, thus reducing complexity. There are two kinds of regularization: L1 and L2. Both add a certain amount to the loss function produced by a layer, multiplied by some regularization parameter λ . Each calculates the additional loss differently. If the weights of the layer are given by the vector $\mathbf{w} = \langle w_1, w_2, \dots, w_n \rangle$, then

$$L1 = \lambda \sum_{i=1}^n |w_i| \quad (\text{B.10})$$

$$L2 = \lambda \sum_{i=1}^n w_i^2. \quad (\text{B.11})$$

When looking at the partial derivatives of these regularization techniques with respect to the i th weight, we get the following results:

$$\frac{\partial L1}{\partial w_i} = \begin{cases} \lambda & \text{if } w_i > 0 \\ -\lambda & \text{if } w_i < 0 \end{cases} \quad (\text{B.12})$$

$$\frac{\partial L2}{\partial w_i} = 2\lambda w_i \quad (\text{B.13})$$

Remembering that gradient descent uses the negative of these partial derivatives to determine how far to adjust the parameters, L1 regularization has the effect of constantly pushing the values of the weights towards 0, while L2 regularization pushes




Figure B.7: Graph showing the difference between a complex function and a simple one. The blue function is more complex than the green one. Image by Wikimedia user Nicoguardo. Licensed under the Creative Commons Attribution 4.0 International license.

against larger parameters. Both have the effect of helping to prevent the weights from becoming too large.

Batch Normalization

Another technique for mitigating overfitting (and generally improving overall performance) is batch normalization. The core idea behind this technique is that a node's value matters less than its variation from the mean. Typically, the elements of a training set are processed by the neural network in batches. Batch normalization leverages this and processes each node's variance among other instances of the same node within a batch.

If the node whose value we are normalizing is in a batch of size n , then its mean and standard deviation can be calculated using the following equations:

$$\mu = \frac{1}{n} \sum_{i=1}^n x_i \quad (\text{B.14})$$

$$\sigma^2 = \frac{1}{n} \sum_{i=1}^n (x_i - \mu)^2. \quad (\text{B.15})$$

The value's normalized value can then be processed using the following equation:

$$\hat{x}_i = \frac{x_i - \mu}{\sqrt{\sigma^2 + \epsilon}} \quad (\text{B.16})$$

The ϵ variable is a small value to ensure that there is no division by 0.

Performance Metrics

To evaluate the performance of each model, the following metrics were used:

Name	Calculation
Accuracy	$\frac{TP+TN}{TP+TN+FP+FN}$
Precision	$\frac{TP}{TP+FP}$
Recall	$\frac{TP}{TP+FN}$
F1	$\frac{2*Precision*Recall}{Precision+Recall}$
MCC	$\frac{TP*TN-FP*FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$

Table B.1: Performance metrics





Figure B.8: Graphic showing the process of cross validation. Image by Wikimedia user Gufosowa. Licensed under the Creative Commons Attribution 4.0 International license.

Cross Validation


To determine the optimal parameters of the neural networks, we used a technique called cross validation. In this process, we take the training set and split it into a given number of folds (if there are 5 folds, it is called 5-fold cross validation, if it has 10 folds, it is 10-fold cross validation, etc.). At each iteration, one fold is used as a validation set, and the others are used as a training set. Each iteration uses a different fold as the validation set, so that each element in the dataset is used in validation. The process is shown in figure B.8.

APPENDIX C


PROTEIN LENGTHS FOR ALL DATASETS




(a) Lengths of proteins for mouse-ear cresses




(b) Lengths of proteins for fruit flies




(c) Lengths of proteins for zebrafish



(d) Lengths of proteins for African clawed frogs




(e) Lengths of proteins for humans




(f) Lengths of proteins for mice

Figure C.1: A histogram comparison of the sequence lengths in each unfiltered dataset.



(g) Lengths of proteins for rats



(h) Lengths of proteins for slime molds





(i) Lengths of proteins for *E. coli*(j) Lengths of proteins for *H. pylori*(k) Lengths of proteins for *P. putida*(l) Lengths of proteins for *B. subtilis*

Figure C.1: A histogram comparison of the sequence lengths in each unfiltered dataset. (cont'd)






(m) Lengths of proteins for *S. solfataricus*(n) Lengths of proteins for *S. pombe*(o) Lengths of proteins for *S. cerevisiae*


Figure C.1: A histogram comparison of the sequence lengths in each unfiltered dataset. (cont'd)




(a) Filtered lengths of proteins for mouse-ear cresses




(b) Filtered lengths of proteins for fruit flies




(c) Filtered lengths of proteins for zebrafish



(d) Filtered lengths of proteins for African clawed frogs




(e) Filtered lengths of proteins for humans




(f) Filtered lengths of proteins for mice

Figure C.2: A histogram comparison of the sequence lengths in each filtered dataset.



(g) Filtered lengths of proteins for rats



(h) Filtered lengths of proteins for slime molds








(i) Filtered lengths of proteins for *E. coli*(j) Filtered lengths of proteins for *H. pylori*(k) Filtered lengths of proteins for *P. putida*(l) Filtered lengths of proteins for *B. subtilis*

Figure C.2: A histogram comparison of the sequence lengths in each filtered dataset.
(cont'd)



(m) Filtered lengths of proteins for *S. solfataricus* (n) Filtered lengths of proteins for *S. pombe*



(o) Filtered lengths of proteins for *S. cerevisiae*

Figure C.2: A histogram comparison of the sequence lengths in each filtered dataset.
(cont'd)