Vulnerability Detection in C/C++ Code with Deep Learning

Zhen Huang

School of Computing, DePaul University, Chicago, IL, USA E-mail: zhen.huang@depaul.edu

Amy Aumpansub

School of Computing, DePaul University, Chicago, IL, USA E-mail: amy.aumpansub@gmail.com

Abstract: Deep learning has been shown to be a promising tool in detecting software vulnerabilities. In this work, we train neural networks with program slices extracted from the source code of C/C++ programs to detect software vulnerabilities. The program slices capture the syntax and semantic characteristics of vulnerability-related program constructs, including API function call, array usage, pointer usage, and arithmetic expression. To achieve a strong prediction model for both vulnerable code and non-vulnerable code, we compare different types of training data, different optimizers, and different types of neural networks. Our result shows that combining different types of characteristics of source code and using a balanced number of vulnerable program slices and non-vulnerable code and non-vulnerable code and non-vulnerable code and non-vulnerable code and non-vulnerable program slices produce a balanced accuracy in predicting both vulnerable code and non-vulnerable code. Among different neural networks, BGRU with the ADAM optimizer performs the best in detecting software vulnerabilities with an accuracy of 92.49%.

Keywords: software vulnerabilities; vulnerability detection; deep learning; neural networks; program analysis.

1 Introduction

Software vulnerabilities pose a significant threat to the security of networks and information. Hackers and malware often take advantage of these vulnerabilities to compromise computer systems, because vulnerabilities enable them to dramatically increase the magnitude and speed of cyber attacks. To incentivize individuals to find such vulnerabilities, renowned software vendors are known to offer rewards as high as \$1 million (Intel Corporation 2020, Microsoft Corporation 2020, Apple Inc. 2020, Facebook 2020).

As manually finding vulnerabilities typically incur considerable effort and time, numerous studies have been dedicated to automatically identify vulnerabilities (Kim et al. 2017, Li et al. 2016, Grieco et al. 2016*a*, Neuhaus et al. 2007*a*, Yamaguchi et al. 2013, 2012*a*). Primarily, these approaches rely on code similarity detection or pattern matching techniques. However, code similarity detection may not effectively identify vulnerabilities that do not result from code duplication, and pattern matching necessitates the expertise of human professionals to define vulnerability patterns.

To address the limitation, neural networks have recently been employed for vulnerability detection (Yang et al. 2015, Shin et al. 2015, White et al. 2016, Wang et al. 2016, Li et al. 2017, Guo et al. 2017, Li, Zou, Xu, Ou, Jin, Wang, Deng & Zhong 2018, Zhou et al. 2019, Lin et al. 2020). Neural networks have gained widespread recognition in fields such as image processing and speech recognition, owing to their ability to deliver highly accurate predictions with minimal dependence on human experts for feature extraction. Given the diverse causes of software vulnerabilities, neural networks automatically extract features and thus mitigate the impact of human bias in feature extraction.

This paper presents our work on using neural networks to create predictive models for automatically detecting vulnerabilities. It consists of four major steps: 1) extracting code relevant to vulnerabilities, 2) converting the extracted code into numeric vectors, 3) training and optimizing neural networks using the numeric vectors, and 4) detecting vulnerabilities using the models generated by the neural networks.

First, we use program slicing to extract syntax and semantic information of four different types of program constructs relevant to vulnerabilities from the source code of target programs. The program constructs include library or API function call, array usage, pointer usage, and arithmetic expression. Each program slice contains the vulnerability-related program construct, and the program statements on which the program construct are control dependent or data dependent.

Second, the extracted program slices are then converted into numeric vectors using the Word2Vector model. Each slice is split into tokens, and the tokens are used as the input to the Word2Vector model, which learns word embedding and outputs the word embedding numeric vectors for the tokens.

Third, the numeric vectors representing the tokens are pre-processed and fed into neural networks. The purpose of the pre-processing is to improve the accuracy of the models generated by the neural networks in vulnerability detection. The pre-processing consists of dataset balancing and integration.

We perform dataset balancing because our dataset has substantially more non-vulnerable program slices than vulnerable program slices, reflecting the fact that programs have much more non-vulnerable code than vulnerable code. To balance the dataset, we downsize the numeric vectors for non-vulnerable program slices.

While prior work (Li, Zou, Xu, Jin, Zhu & Chen 2018) trains individual models on each type of vulnerability-related program construct, our work trains on the data integrated from all types of program constructs. Our results demonstrate that the model built from the integrated dataset outperforms the individual models created from separate datasets.

To create a robust predictive model, we fine-tune the neural networks by experimenting with various hyperparameters, including optimizers and gating mechanisms, during model development. Our experiments show that the ADAM optimizer and bidirectional RNNs achieve the best results.

Lastly, we use the trained models to identify vulnerabilities in our dataset. Our BGRU model outperforms the BLSTM model. It achieves an accuracy rate of 94.6% on the training set and 92.4% on the test set.

The major contributions of this paper is as follows:

- We show that the accuracy of the model built on the combined dataset surpasses the models built on individual dataset.
- By balancing the ratio of vulnerable data points (class 1) and non-vulnerable data points (class 0), the model performs well with a high balanced accuracy rate of 93% which is comparable to that of a training set. The high sensitivity and specificity imply the model has a good ability in explaining both vulnerability and non-vulnerability classes.
- We compare different types of neural networks and show that BGRU performs the best. The model built with BGRU achieves an accuracy rate of 94.89% by utilizing 10X more data points.
- We implement a chain of tools for generating the model from program slices and open source the tools at https://gitlab.com/vulnerability_analysis/vulnerability_detection/.

The paper is structured into six sections. Section 2 presents information on the dataset. Section 3 describes the details on fine-tuning the models. Section 4 shows evaluation results. Section 5 discusses related work. Finally, we conclude in Section 6. This paper expands upon the ideas presented in Aumpansub & Huang (2021).

2 Dataset

Our work uses the dataset of C/C++ programs collected by Li, Zou, Xu, Jin, Zhu & Chen (2018). The dataset includes 1,592 programs from the National Vulnerability Database (NVD) and 14,000 programs from the Software Assurance Reference Dataset (SARD). These programs were pre-processed and transformed to 420,627 program slices called semantic vulnerability candidates (SeVC) which contain 56,395 vulnerable slices (13.5 % of program slices) and 364,232 non-vulnerable slices (86.5 % of program slices). The program slices are then transformed into numeric vectors that will be used as inputs to neural networks.

The program slices were created by extracting statements relevant to four types of vulnerability-relevant program constructs:

- Library or API Function Call (API). This type of program slices is associated with library or API functions calls for 811 C/C++ library/API function calls. This type represents 15.3% of total slices, comprising 13,603 vulnerable slices and 50,800 non-vulnerable slices.
- Array Usage (AU). This type of program slices is related to the use of arrays such as array element access, accounting for 10% of total slices which contain 10,926 vulnerable slices and 31,303 non-vulnerable slices.

- **Pointer Usage (PU).** This type of program slices is related to the use of pointer arithmetic and dereferences. This type represents 69.4% of total slices which include 28,391 vulnerable slices and 263,450 non-vulnerable slices.
- Arithmetic Expression (AE). This type of program slices is associated with arithmetic expressions such as integer additions and subtractions, which represents 5.3% of total slices, comprising 3,475 vulnerable slices and 18,679 non-vulnerable slices.

2.1 Generating Program Slices

The program slices are generated in two phases. First, syntax-based vulnerability candidates (SyVCs) are extracted from programs, based on the abstract syntax trees (ASTs) of the programs. Each SyVC encapsulates the syntax characteristics of a vulnerability-related program construct. Second, semantics-based vulnerability candidates (SeVCs) are generated from SyVCs by generating a program dependency graph (PDG) for each function of the programs and extending each SyVC with data dependency and control dependency information from PDGs. Each SeVC is a program slice that contains semantic and syntax information related with a vulnerability-related program construct. We define the type of a program slice as the type of program construct on which the program slice is generated.

The process of generating program slices is illustrated in Figure 1. The Joern package in Python was used to parse the source code and generate PDG. More details on program slice generation can be found in Li, Zou, Xu, Jin, Zhu & Chen (2018).



Figure 1 Generating Program Slices fro Source Code.

2.2 Transforming Program Slices into Vectors

To use the program slices with neural networks, they need to be transformed into numeric vectors. Each slice is first split into a list of tokens in which all comments and white spaces were removed. It is also mapped to the list of relevant functions.

The list of tokens for each slice are stored in a pickle file and labeled with a unique ID. Each pickle file contains five elements: a list of tokens, a target label (0/1), a list of functions, vulnerability type, and the ID of the slice. A target label of 0 indicates that the slice is non-vulnerable, while a target label of 1 indicates that the slice is vulnerable.

The list of tokens from each pickle file is converted intto vectors using the Word2Vector model, which converts tokens to vectors based on cosine similarity distance, measuring the angle between vectors. A higher similarity score indicates a higher similarity and a closer distance between tokens (Mikolov et al. 2013). The cosine similarity is computed as follows:

For each program slice, the output of the Word2Vector model is a $30 \times n$ array, where 30 is the dimension of the columns and n is the dimension of the rows. Each row is the word embedding for one token and thus n is the number of tokens in the program slice.

$$sim(X,Y) = \frac{X \bullet Y}{\|X\| \times \|y\|} = \frac{\sum_{i} (x_i \times y_i)}{\sqrt{\sum_{i} x_i^2} \times \sqrt{\sum_{i} y_i^2}}$$

Figure 2 Cosine similarity.

The visualization of tokens in the Word2Vector model is shown in Figure 3. As we can see, different program slice types have substantially different distributions of cosine similarities. This indicate that different program slice types convey different characteristics of vulnerabilities.



Figure 3 Visualized tokens in W2V model for each program slice type.

3 Model Optimization

In this section, we describe the steps that we took to find optimal pre-processing techniques and neural network models. We use a subset of the dataset for the majority of our experiments. The subset includes randomly chosen 30,000 vector arrays from the total 420,627 vector arrays, each of which corresponds to a program slice. The subset is split into a training set of 24,000 vector arrays and a testing set of 6,000 vector arrays. First, we compare the results on individual program slice types and the results on combined program slice types. Second, we show the results using an imbalanced dataset and the results using a balanced dataset. Third, we experiment with various optimizers including SGD, ADAMAX, and ADAM. Last, we discuss and compare the results using different RNNs.

3.1 Combining Program Slice Types

From the visualization of Word2Vector models for each program slice type, as shown in Figure 3, we note that different program slice types capture different characteristics of vulnerabilities, so we explore the use of a dataset combined from all different program slice types.

We perform a preliminary study using 1,000 randomly chosen program slices from each individual program slice types, collectively called individual datasets, and 1,000 randomly chosen program slices from all different program slice types, called combined dataset. We compare the accuracy, sensitivity, and specificity for the models built using individual datasets and the model built using the combined dataset. Table 1 shows our result.

Туре	Accuracy	Sensitivity	Specificity
API	53%	69%	46%
AU	64%	79%	62%
PU	38%	83%	31%
AE	61%	61%	62%
COMBINED	61%	91%	53%

 Table 1
 Comparison between individual datasets and combined dataset.

The model built using the combined dataset, i.e. combined model, outperforms the models built using individual datasets, i.e. individual models, in detecting the target class 1 (vulnerable) code, as the sensitivity of the model is 91%, the highest among the all models. In detecting the target class 0 (non-vulnerable) code, the combined model performs considerably better than the API model and PU model, while it performs slightly worse than the AU model and AE model. As a result, we consider the combined dataset more appropriate for predicting vulnerabilities.

3.2 Balancing Dataset

In general, the dataset used for training should have a balanced number of class 0 samples (non-vulnerable code) and class 1 samples (vulnerable code) to ensure that the model can produce unbiased predictions. However, Li, Zou, Xu, Jin, Zhu & Chen (2018) used an imbalanced dataset in which class 1 samples only account for 15.6% of the total program slices while class 0 samples account for 84.4% of the total program slices.

To illustrate the issue, we compute the confusion matrix for the model built with imbalanced dataset (75% of class 0 and 25% of class 1). As presented in Table 2, the model has considerably higher accuracy in predicting class 0 samples, as its specificity and negative prediction are remarkably higher than its sensitivity and precision, respectively. Its accuracy rate is biased towards class 0.

Predicted Class				
	Positive	Negative	Rate	
Positive	8.0	48.0	0.14285	Sensitivity
Negative	14.0	130.0	0.90277	Specificity
	0.36363	0.73033	0.60999	Accuracy
	Precision	Negprediction		

 Table 2
 Confusion matrix for imbalanced dataset.

In order to address the issue, we re-sample the training set using a down-sampling method, which randomly removes samples from the majority class (label 0) of the training

set to make the number of class 0 samples the same as the number of class 1 samples. The new training set has a balanced samples, containing 50% vulnerable samples and 50% non-vulnerable samples. Figure 4 shows the process of down-sampling.

Because neural networks require all vector input to have the same dimension, we also adjust our vector arrays to have the same number of rows, i.e. same vector lengths. We compute the average number of rows of the vector arrays, and use it as the threshold to adjust the vector lengths. If a vector array has a vector length less than the average length, we append the vector array with zero vectors. If a vector array has a vector length larger than the average length, we truncate the vector array.



Figure 4 Down-sampling and vector adjustment.

As shown in Table 3, the model built using the balanced dataset has approximately the same accuracy in predicting class 0 samples and class 1 samples. This shows that balancing the dataset is critical for the model to have balanced prediction power for both classes.

Predicted Class				
	Positive	Negative	Rate	
Positive	1186.0	167.0	0.87916	Sensitivity
Negative	672.0	4018.0	0.86408	Specificity
	0.65236	0.96108	0.86747	Accuracy
	Precision	Negprediction		

 Table 3
 Confusion matrix for balanced dataset.

3.3 Selecting Optimizers

Different optimizers can be applied to optimize neural networks. They are algorithms for finding the optimal parameters for a model during the training process by adjusting the weights and biases in the model iteratively until they converge on a minimum loss value. Some of the most popular optimizers include SGD, Momentum, ADAMGRAD, RMS Prop, ADAM and ADAMAX. In order to find the best optimizer for our neural networks, we explore three different optimizers: SGD, ADAM, and ADAMAX.

SGD computes the gradient of the loss function based on a randomly chosen subset of the training data instead of the entire training data. Comparing to the standard gradient descent, SGD can converge faster and use less memory storage.

ADAM computes individual learning rates for different parameters. It keeps track of a changing average of the gradient's first and second moments, which are respectively

 Table 4
 Accuracy rate with different optimizers.

Туре	ADAMAX	SGD	ADAM
API	86.7%	63.1%	89.5%
AU	86.0%	58.6%	89.2%
PU	82.4%	62.3%	90.9%
AE	83.1%	67.1%	90.5%

the mean and variance of the gradients. ADAM is appropriate for large dataset and/or parameters, with non-stationary objectives, and for problems with very noisy and/or sparse gradients Kingma & Ba (2015).

ADAMAX is a variant of ADAM. Similar to ADAM, ADAMAX also keeps track of a changing average of the gradient's mean and variance of the gradients. Different from ADAM, ADAMAX uses the L-infinity norm of the gradients instead of the second moment of the gradients. ADAMAX is appropriate for the scenarios in which the gradients are sparse or have a high variance.

The accuracy of the models using ADAMAX, SGD, and ADAM optimizer is presented in Table 4. We can see that ADAM performs the best among the three optimizers for all program slice types. ADAM achieves an average accuracy rate of 90.0%. This is approximately 5% higher than the accuracy rate of ADAMAX, which is used in prior work (Li, Zou, Xu, Jin, Zhu & Chen 2018). As a result, we choose to use ADAM for our neural networks.

3.4 Comparing RNNs

In this section, we discuss and compare the performance of different neural networks. First, we discuss GRU and LSTM. Second, we compare LSTM with BLSTM. Last, we analyze BGRU and BLSTM.

GRU vs. LSTM. Comparing to LSTM, GRU has no explicit memory unit, no forget gate and update gate. GRU also has fewer number of hyperparameters. With a simpler architecture, GRU trains faster than LSTM. However, GRU may have lower accuracy rate than LSTM, because LSTM comprises both update gate and forget gate and remembers longer sequences than GRU, although LSTM is comparable to GRU on sequence modeling.

BLSTM vs. LSTM. A bidirectional recurrent neural network (RNN) has two layers sideby-side. It provides the original input sequence to the first layer and a reversed copy of the input sequence to the second layer. Bidirectional RNNs are found to be more effective than regular RNNs, because it can overcome the limitations of a regular RNN (Schuster & Paliwal 1997). A regular RNN preserves only information of the past, while a bidirectional RNN has access to the past information as well as the future information. Therefore the output of a bidirectional RNN is generated from both the past context and future context, and that leads to a better prediction and classifying capability. Our experiment in training LSTM and BLSTM models on a subset of our dataset also indicates that BLSTM outperforms LSTM using the same hyperparameters, as shown in Figure 5.

The result of this experiment shows that the BLSTM model has a lower loss rate of 0.58, as compared to the LSTM model's loss rate of 0.60. The BLSTM model also has a higher accuracy rate of 64.2% than the LSTM model's accuracy rate of 62.8%. Note that in this experiment both models were fit with the same input parameters on a small dataset, which includes 1,000 program slices, so the accuracy rates are not high.

BLS	ТΜ
-----	----

LSTM

start	start
Enoch 1/10	Enoch 1/10
	Eboci 1/10
9/9 [===================================	9/9 [===================================
Epoch 2/10	Epoch 2/10
9/9 [======================] - 11s 1s/step - loss: 0.7036 - accuracy: 0.4549	9/9 [==================] - 5s 591ms/step - loss: 0.6990 - accuracy: 0.4792
Epoch 3/10	Epoch 3/10
9/9 [======================] - 13s 1s/step - loss: 0.6761 - accuracy: 0.6076	9/9 [==================] - 5s 537ms/step - loss: 0.6852 - accuracy: 0.5903
Epoch 4/10	Epoch 4/10
9/9 [=======================] - 13s 1s/step - loss: 0.6690 - accuracy: 0.5938	9/9 [=======================] - 5s 540ms/step - loss: 0.6764 - accuracy: 0.5938
Epoch 5/10	Epoch 5/10
9/9 [============================] - 13s 1s/step - loss: 0.6559 - accuracy: 0.6111	9/9 [=================================] - 5s 533ms/step - loss: 0.6546 - accuracy: 0.6250
Epoch 6/10	Epoch 6/10
9/9 [===================================	9/9 [=======================] - 5s 534ms/step - loss: 0.6443 - accuracy: 0.5938
Epoch 7/10	Epoch 7/10
9/9 [===================================	9/9 [=======================] - 5s 531ms/step - loss: 0.6478 - accuracy: 0.6007
Enoch 8/10	Epoch 8/10
9/9 [] - 11s 1s/step - loss: 0.6092 - accuracy: 0.6632	9/9 [===================================
Enoch 9/10	Epoch 9/10
9/9 [===================================	9/9 [===============================] - 5\$ 533ms/step - Loss: 0.6364 - accuracy: 0.5833
Ench 10/10	Ebocu 10/10
	9/9 [================================] - 5s 572ms/step - loss: 0.6096 - accuracy: 0.6285
9/9 [===================================	

Figure 5 Model fitting of BLSTM and LSTM.

Predicted Class				
	Positive	Negative	Rate	
Positive	185.0	41.0	0.81858	Sensitivity
Negative	295.0	478.0	0.61837	Specificity
	0.38542	0.92100	0.66366	Accuracy
	Precision	Negprediction		

Table 5Confusion matrix for BGRU with 5,000 samples.

Table 6Confusion matrix for BLSTM with 5,000 samples.

Predicted Class				
	Positive	Negative	Rate	
Positive	185.0	41.0	0.81858	Sensitivity
Negative	321.0	452.0	0.58473	Specificity
	0.36561	0.91684	0.63764	Accuracy
	Precision	Negprediction		

BGRU vs. BLSTM. Table 5 and Table 6 show the confusions matrices for BGRU and BLSTM respectively. The models were trained on 4,000 samples and tested on 1,000 samples. The decision threshold is set to 0.5 for validation. The BGRU model outperforms the BLSTM in most metrics except for the sensitivity. It has a higher accuracy, precision, and specificity, which indicates its stronger capability to predict both vulnerable code and non-vulnerable code.

Table 7Confusion matrix for BGRU with 30,000 samples.

Predicted Class				
	Positive	Negative	Rate	
Positive	731.0	77.0	0.90470	Sensitivity
Negative	1022.0	4170.0	0.80316	Specificity
	0.41699	0.98187	0.81683	Accuracy
	Precision	Negprediction		

Table 8 Confusion matrix for BLSTM with 30,000 samples.

Predicted Class				
	Positive	Negative	Rate	
Positive	663.0	145.0	0.82054	Sensitivity
Negative	1095.0	4097.0	0.78910	Specificity
	0.37713	0.96582	0.79333	Accuracy
	Precision	Negprediction		

BGRU outperforms BLSTM in all the metrics on a larger dataset of 30,000 samples. As shown in Table 7 and Table 8, the BGRU model has a sensitivity of 90%, which is 8% higher than that of the BLSTM model. This indicates that the BGRU model can predict the vulnerable code better than the BLSTM model. Table 9 and Table 10 show that the BGRU model also performs better than the BLSTM model on a dataset of 100,000 samples. Comparing to the BLSTM model, the BGRU model has 3% higher accuracy and specificity, and approximately the same sensitivity.

Table 9Confusion matrix for BGRU with 100,000 samples.

Predicted Class				
	Positive	Negative	Rate	
Positive	2383.0	287.0	0.89251	Sensitivity
Negative	1506.0	15824.0	0.91310	Specificity
	0.61275	0.98219	0.91035	Accuracy
	Precision	Negprediction		

Table 10 Confusion matrix for BLSTM with 100,000 samples.

Predicted Class				
	Positive	Negative	Rate	
Positive	2408.0	262.0	0.90187	Sensitivity
Negative	2193.0	15137.0	0.87346	Specificity
	0.52336	0.98300	0.87725	Accuracy
	Precision	Negprediction		

4 Evaluation

In this section, we evaluate the accuracy of our neural networks in predicting vulnerable code and non-vulnerable code. We first show the results on individual program slice types, then show the results on the combined program slice types. We build the models using BGRU for all the evaluations.

4.1 Individual Program Slice Types

The dataset contains program slices created for four types of vulnerability-related program constructs: library or API functions (API), array usage (AU), pointer usage (PU), and arithmetic expressions (AE). We build individual models for each program slice type. We use a dataset of 6,000 program slices for this evaluation. Our focus is to find a threshold on the prediction results that will have the best accuracy.

Table 11 presents the threshold for the BGRU model to achieve the highest accuracy rate and F1 score for each type of program slices. As we can, the threshold to get the highest accuracy rate ranges from 0.4 to 0.65, with a mean of 0.5, and the threshold to get the highest F1 score ranges from 0.55 to 0.7, with a mean of 0.625.

Туре	Threshold for Accuracy	Threshold for F1
API	0.55	0.65
AU	0.4	0.55
PU	0.4	0.6
AE	0.65	0.7

 Table 11
 Prediction thresholds for different program slice types.

API. The effects of different thresholds for API slices are shown in Table 12. A prediction threshold of 0.55 achieves the highest accuracy of 0.872 while a predication threshold of 0.65 achieves the highest F1 score of 0.759.

Threshold	Recall	Precision	Specificity	F1	Accuracy
Thieshold	Recall	TICCISION	specificity	1.1	Accuracy
0.4	0.906597	0.617677	0.837204	0.734755	0.871901
0.45	0.889548	0.630252	0.848602	0.737781	0.869075
0.5	0.87917	0.652365	0.864086	0.748974	0.871628
0.55	0.870274	0.668184	0.874624	0.755956	0.872449
0.6	0.858414	0.678781	0.882151	0.758101	0.870282
0.65	0.841464	0.690809	0.890753	0.75869	0.866058
0.7	0.811712	0.702824	0.90043	0.753354	0.856071

Table 12Prediction threshold for API slices.

Array Usage. The effects of different thresholds for AU slices are shown in Table 13, a prediction threshold of 0.4 achieves the highest accuracy of 0.878 while a predication threshold of 0.55 achieves the highest F1 score of 0.812.

Pointer Usage. The effects of different thresholds for PU slices are shown in Table 14, a prediction threshold of 0.4 achieves the highest accuracy of 0.837 while a predication threshold of 0.6 achieves the highest F1 score of 0.693.

Arithmetic Expression. The effects of different thresholds for AE slices are shown in Table 15, a prediction threshold of 0.65 achieves the highest accuracy of 0.878 while a predication threshold of 0.7 achieves the highest F1 score of 0.677.

Table 13Prediction threshold for Array Usage slices.

Threshold	Recall	Precision	Specificity	F1	Accuracy
0.4	0.946099	0.704511	0.809183	0.807625	0.877641
0.45	0.931725	0.713724	0.820291	0.808283	0.876008
0.5	0.918891	0.724696	0.83214	0.810321	0.856031
0.55	0.904517	0.737238	0.844977	0.812356	0.874747
0.6	0.86961	0.745599	0.857319	0.802844	0.86131
0.65	0.826489	0.770704	0.881758	0.797622	0.854123
0.7	0.766427	0.800966	0.908418	0.783316	0.837422

 Table 14
 Prediction threshold for Pointer Usage slices.

Threshold	Recall	Precision	Specificity	F1	Accuracy
0.4	0.885281	0.556715	0.788207	0.683565	0.836744
0.45	0.872294	0.568139	0.80078	0.688105	0.836537
0.5	0.844156	0.582669	0.818339	0.689452	0.831248
0.55	0.822511	0.599054	0.834598	0.69322	0.828554
0.6	0.798701	0.612618	0.848255	0.693392	0.823478
0.65	0.771284	0.624051	0.860395	0.6899	0.815839
0.7	0.731602	0.641366	0.877086	0.683519	0.804344

 Table 15
 Prediction threshold for Arithmetic Expression slices.

Threshold	Recall	Precision	Specificity	F1	Accuracy
0.4	0.936324	0.444979	0.784167	0.603263	0.860246
0.45	0.930535	0.46259	0.800214	0.617972	0.865375
0.5	0.927641	0.479073	0.813587	0.631838	0.870614
0.55	0.920405	0.496487	0.827494	0.64503	0.87395
0.6	0.914616	0.514239	0.840332	0.658333	0.877474
0.65	0.901592	0.53339	0.854239	0.670253	0.877915
0.7	0.885673	0.548387	0.865205	0.677366	0.875439

4.2 Combined Program Slice Types

We combine the total 420,067 programs slices into one dataset, comprising 64,403, 42,229, 291,281, and 22,154 from API, AU, PU, and AE types, respectively. The combined dataset is split into a training set and a test set with the 80:20 ratio. The training set is then down-sampled to ensure that the target classes (vulnerable and non-vulnerable) in it are balanced.

Our BGRU model is built with the ADAM optimizer. The hyperparameters of the model include 256 neuron units with 2 hidden layers. The Tanh function is applied to produce the outputs of 2 hidden layers and the Sigmoid function is applied to compute activation outputs in the last layer. The learning rate is 0.1 with a batch size of 32. The binary cross-entropy loss function is used as it can speed up the convergence.

We illustrates the learning process in Figure 6. The learning process is faster in the beginning, as the loss rate significantly decreases in epoch 1 to 3. The accuracy rate increases as the training process goes from epoch 1 to 10. The model has the highest accuracy rate of 94.89% in epoch 9 and starts to decrease in epoch 10 as the error rate is no longer reduced.

The output of the model ranges between 0 and 1, as the Sigmoid function is applied to the output layer.

Epoch 1/10										
2824/2824 [==		10	13285s	5s/step	-	loss:	0.4097	æ	accuracy:	0.8209
Epoch 2/10										
2824/2824 [==]	17	14041s	5s/step	-	loss:	0.2432	-	accuracy:	0.9039
Epoch 3/10										
2824/2824 [==]) a	14114s	5s/step	-	loss:	0.1904	-	accuracy:	0.9267
Epoch 4/10										
2824/2824 [==]	-	14219s	5s/step	-	loss:	0.1652	2	accuracy:	0.9376
Epoch 5/10										
2824/2824 [==		5	14375s	5s/step	-	loss:	0.1532		accuracy:	0.9419
Epoch 6/10										
2824/2824 [==		ie.	14330s	5s/step	-	loss:	0.1444	-	accuracy:	0.9459
Epoch 7/10										
2824/2824 [==	[]	-	14374s	5s/step	-	loss:	0.1398	-	accuracy:	0.9472
Epoch 8/10										
2824/2824 [==			14512s	5s/step	-	loss:	0.1373	2	accuracy:	0.9480
Epoch 9/10										
2824/2824 [==]	12	14566s	5s/step	-	loss:	0.1371	-	accuracy:	0.9489
Epoch 10/10										
2824/2824 [==]	5	14642s	5s/step	-	loss:	0.1371	(\Box)	accuracy:	0.9482



Predicted Class				
	Positive	Negative	Rate	
Positive	10768.0	439.0	0.96082	Sensitivity
Negative	5898.0	67019.0	0.91911	Specificity
	0.64610	0.99349	0.92467	Accuracy
	Precision	Negprediction		

 Table 16
 Confusion matrix for test set.

Table 16 shows the confusion matrix for the test set. We can see that the model performs well in predicting both target classes (vulnerable code and non-vulnerable code), as both the sensitivity and specificity are over 90%.

As presented in Figure 7, the F1 score increases and the balanced accuracy decreases while the threshold increases. The peak point of the balanced accuracy is achieved when the threshold is 0.5. The peak point of the F1 score is achieved when the threshold is 0.8.

Overall, the model fitted with the combined dataset performs well with a high accuracy rate of 92.5%. Its high sensitivity and specificity indicates that it has a good capability in predicting both vulnerable code and non-vulnerable code, although the model performs better in predicting non-vulnerable code than vulnerable code, as it has a negative prediction rate of 99.3%.

5 Related Work

Many approaches have been proposed to detect and address vulnerabilities (Valeur et al. 2005, Neuhaus et al. 2007*b*, Dessiatnikoff et al. 2011, Shin et al. 2011, Yamaguchi et al.



Figure 7 F1 v.s. Accuracy Rate for Different Thresholds.

2012*b*, Zheng & Zhang 2013, Huang & Lie 2014, Grieco et al. 2016*a*, Li et al. 2016, Wu et al. 2017, Huang & Lie 2017, David et al. 2018, Li, Zou, Xu, Ou, Jin, Wang, Deng & Zhong 2018, Li, Zou, Xu, Jin, Zhu & Chen 2018, Chernis & Verma 2018, Wang et al. 2010, Huang & Tan 2019, Li et al. 2019, Lin et al. 2020, Zagane et al. 2020, Li, Zou, Xu, Jin, Zhu & Chen 2021, Huang & Yu 2021, Li, Wang & Nguyen 2021, Huang et al. 2021, Eshghie et al. 2021, Hin et al. 2022, Huang & White 2022, Fu & Tantithamthavorn 2022, Aumpansub & Huang 2022, Cao et al. 2022). They can be broadly categorized as rule-based approaches and learning-based approaches.

Rule-based approaches detect the existence of vulnerabilities using predefined rules, which typically characterize vulnerable and non-vulnerable program code structures (David et al. 2018) or behaviors (Wang et al. 2010). Rule-based approaches follow the predefined rules to analyze program code or program behaviors. These analyses can be performed dynamically (Huang & Yu 2021), which execute target programs, or statically (Zheng & Zhang 2013), which examine target programs without executing them. Rule-based approaches identify a vulnerability when a predefined rule finds a match of vulnerable code structures or behaviors. A major disadvantage of rule-based approaches is that the predefined rules require considerable manual effort and time to generate.

As learning-based approaches have thrived in a myriad of areas, particularly in software security and reliability (Mickens et al. 2007, Yuan et al. 2011, Huang & Lie 2014, Grieco et al. 2016b, Wang et al. 2016, Long & Rinard 2016, Huang & Lie 2017, Cummins et al. 2018, Li et al. 2019, Tien et al. 2020), they have also been leveraged in vulnerability detection. Learning-based approaches extract characteristics of program code or behaviors automatically and identify vulnerabilities based on these characteristics. Conventional machine learning approaches learn characteristics of vulnerabilities using various human-defined features such as source code text features in the source code (Chernis & Verma

2018), complexity, code churn, and developer activity metrics (Shin et al. 2011), abstract syntax trees (Yamaguchi et al. 2012*b*), function imports and function calls (Neuhaus et al. 2007*b*). Similar to rule-based approaches, a main drawback of conventional machine learning approaches is that they require considerable human effort to define these features.

Recent approaches use deep learning on program code to detect vulnerabilities so that no human experts is needed to define features (Li, Zou, Xu, Ou, Jin, Wang, Deng & Zhong 2018, Li, Zou, Xu, Jin, Zhu & Chen 2018, Li et al. 2019, Zagane et al. 2020, Li, Zou, Xu, Jin, Zhu & Chen 2021, Hin et al. 2022, Li, Wang & Nguyen 2021, Fu & Tantithamthavorn 2022, Aumpansub & Huang 2022, Cao et al. 2022). They typically use neural networks to automatically build classification models from a large number of program samples. Deep learning approaches have been shown to have better accuracy than conventional machine learning approaches (Wu et al. 2017). However, most of them either rely on one type of training data or use imbalanced training data. Our work differs from them by using a balanced dataset that combines different types of training data.

6 Conclusion

We present our work on detecting software vulnerabilities using neural networks. In this work, we train neural networks with program slices extracted from the source code of 15,592 C/C++ programs. The program slices encapsulate characteristics of different types of vulnerability-related program constructs. We compare different types of training data and different types of neural networks. Our results show that the model based on the combined slices of different program construct types outperforms the models based on the slices of individual program construct types. Using a balanced number of vulnerable program slices and non-vulnerable program slices ensures that the model has a balanced accuracy in predicting both vulnerable code and non-vulnerable code. We find that BGRU performs the best among other neural networks. It achieves an accuracy of 94.89%, with a sensitivity of 96.08% and a specificity of 91.91%.

References

- Apple Inc. (2020), 'Apple Security Bounty', https://developer.apple.com/ security-bounty/.
- Aumpansub, A. & Huang, Z. (2021), Detecting Software Vulnerabilities Using Neural Networks, *in* 'Proceedings of the 13th International Conference on Machine Learning and Computing', ICMLC 2021, ACM, pp. 166–171. URL: https://doi.org/10.1145/3457682.3457707
- Aumpansub, A. & Huang, Z. (2022), Learning-based Vulnerability Detection in Binary Code, *in* 'Proceedings of the 14th International Conference on Machine Learning and Computing', ICMLC 2022, ACM, pp. 266–271. URL: https://doi.org/10.1145/3529836.3529926
- Cao, S., Sun, X., Bo, L., Wu, R., Li, B. & Tao, C. (2022), Mvd: memory-related vulnerability detection based on flow-sensitive graph neural networks, *in* 'Proceedings of the 44th International Conference on Software Engineering', pp. 1456–1468.

- Chernis, B. & Verma, R. (2018), Machine Learning Methods for Software Vulnerability Detection, *in* 'Proceedings of the Fourth ACM International Workshop on Security and Privacy Analytics', IWSPA '18, ACM, p. 31–39. URL: https://doi-org.ezproxy.depaul.edu/10.1145/3180445.3180453
- Cummins, C., Petoumenos, P., Murray, A. & Leather, H. (2018), 'Compiler fuzzing through deep learning', *ISSTA 2018 Proceedings of the 27th ACM SIGSOFT International Symposium on Software Testing and Analysis* pp. 95–105.
- David, Y., Partush, N. & Yahav, E. (2018), 'Firmup: Precise static detection of common vulnerabilities in firmware', *ACM SIGPLAN Notices* **53**(2), 392–404.
- Dessiatnikoff, A., Akrout, R., Alata, E., Kaâniche, M. & Nicomette, V. (2011), A clustering approach for web vulnerabilities detection, *in* '2011 IEEE 17th Pacific Rim International Symposium on Dependable Computing', IEEE, pp. 194–203.
- Eshghie, M., Artho, C. & Gurov, D. (2021), Dynamic vulnerability detection on smart contracts using machine learning, *in* 'Proceedings of the 25th International Conference on Evaluation and Assessment in Software Engineering', pp. 305–312.

Facebook (2020), 'Facebook', https://www.facebook.com/whitehat/.

- Fu, M. & Tantithamthavorn, C. (2022), Linevul: A transformer-based line-level vulnerability prediction, *in* 'Proceedings of the 19th International Conference on Mining Software Repositories', pp. 608–620.
- Grieco, G., Grinblat, G. L., Uzal, L., Rawat, S., Feist, J. & Mounier, L. (2016a), Toward large-scale vulnerability discovery using machine learning, *in* 'Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy', CODASPY '16, Association for Computing Machinery, New York, NY, USA, p. 85–96. URL: https://doi.org/10.1145/2857705.2857720
- Grieco, G., Grinblat, G. L., Uzal, L., Rawat, S., Feist, J. & Mounier, L. (2016b), Toward large-scale vulnerability discovery using machine learning, *in* 'Proceedings of the Sixth ACM Conference on Data and Application Security and Privacy', CODASPY '16, Association for Computing Machinery, New York, NY, USA, p. 85–96. URL: https://doi.org/10.1145/2857705.2857720
- Guo, J., Cheng, J. & Cleland-Huang, J. (2017), Semantically enhanced software traceability using deep learning techniques, *in* '2017 IEEE/ACM 39th International Conference on Software Engineering (ICSE)', IEEE, pp. 3–14.
- Hin, D., Kan, A., Chen, H. & Babar, M. A. (2022), Linevd: Statement-level vulnerability detection using graph neural networks, *in* 'Proceedings of the 19th International Conference on Mining Software Repositories', pp. 596–607.
- Huang, Z., Jaeger, T. & Tan, G. (2021), Fine-grained Program Partitioning for Security, *in* 'Proceedings of the 14th European Workshop on Systems Security', EuroSec '21, Association for Computing Machinery, New York, NY, USA, pp. 21–26. URL: https://doi.org/10.1145/3447852.3458717

- Huang, Z. & Lie, D. (2014), Ocasta: Clustering Configuration Settings for Error Recovery, *in* 'Proceedings of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks', DSN '14, pp. 479–490. (Acceptance Rate: 30.3%, 56 out of 185).
- Huang, Z. & Lie, D. (2017), 'SAIC: Identifying Configuration Files for System Configuration Management', arXiv:1711.03397.
- Huang, Z. & Tan, G. (2019), Rapid Vulnerability Mitigation with Security Workarounds, *in* 'Proceedings of the 2nd NDSS Workshop on Binary Analysis Research', BAR '19.
- Huang, Z. & White, M. (2022), Semantic-Aware Vulnerability Detection, *in* 'Proceedings of 2022 IEEE International Conference on Cyber Security and Resilience', IEEE, pp. 68–75.
- Huang, Z. & Yu, X. (2021), Integer Overflow Detection with Delayed Runtime Test, *in* 'Proceedings of the 16th International Conference on Availability, Reliability and Security, Vienna, Austria, August 17-20, 2021', ARES 2021, ACM, pp. 28:1–28:6. URL: https://doi.org/10.1145/3465481.3465771
- Intel Corporation (2020), 'Intel Bug Bounty Program', https://www.intel.com/ content/www/us/en/security-center/bug-bounty-program.html.
- Kim, S., Woo, S., Lee, H. & Oh, H. (2017), Vuddy: A scalable approach for vulnerable code clone discovery, *in* '2017 IEEE Symposium on Security and Privacy (SP)', IEEE, pp. 595–614.
- Kingma, D. P. & Ba, J. (2015), Adam: A Method for Stochastic Optimization, *in* 'The 3rd International Conference for Learning Representations'.
- Li, J., He, P., Zhu, J. & Lyu, M. R. (2017), Software defect prediction via convolutional neural network, *in* '2017 IEEE International Conference on Software Quality, Reliability and Security (QRS)', IEEE, pp. 318–328.
- Li, Y., Wang, S. & Nguyen, T. N. (2021), Vulnerability detection with fine-grained interpretations, *in* 'Proceedings of the 29th ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering', pp. 292–303.
- Li, Z., Zou, D., Tang, J., Zhang, Z., Sun, M. & Jin, H. (2019), 'A comparative study of deep learning-based vulnerability detection system', *IEEE Access* 7, 103184–103197.
- Li, Z., Zou, D., Xu, S., Jin, H., Qi, H. & Hu, J. (2016), Vulpecker: An automated vulnerability detection system based on code similarity analysis, *in* 'Proceedings of the 32nd Annual Conference on Computer Security Applications', ACSAC '16, Association for Computing Machinery, New York, NY, USA, p. 201–213.
 URL: https://doi.org/10.1145/2991079.2991102
- Li, Z., Zou, D., Xu, S., Jin, H., Zhu, Y. & Chen, Z. (2018), 'Sysevr: A framework for using deep learning to detect software vulnerabilities'.
- Li, Z., Zou, D., Xu, S., Jin, H., Zhu, Y. & Chen, Z. (2021), 'Sysevr: A framework for using deep learning to detect software vulnerabilities', *IEEE Transactions on Dependable and Secure Computing* **19**(4), 2244–2258.

- Li, Z., Zou, D., Xu, S., Ou, X., Jin, H., Wang, S., Deng, Z. & Zhong, Y. (2018), Vuldeepecker: A deep learning-based system for vulnerability detection, *in* 'Proceedings of the 25th Annual Network and Distributed System Security Symposium', NDSS, pp. 100–116.
- Lin, G., Wen, S., Han, Q.-L., Zhang, J. & Xiang, Y. (2020), 'Software vulnerability detection using deep neural networks: a survey', *Proceedings of the IEEE* **108**(10), 1825–1848.
- Long, F. & Rinard, M. (2016), 'Automatic patch generation by learning correct code', *SIGPLAN Not.* **51**(1), 298–312. **URL:** *https://doi.org/10.1145/2914770.2837617*
- Mickens, J., Szummer, M. & Narayanan, D. (2007), Snitch: interactive decision trees for troubleshooting misconfigurations, *in* 'SYSML'07: Proceedings of the 2nd USENIX workshop on Tackling computer systems problems with machine learning techniques', USENIX Association, Berkeley, CA, USA, pp. 1–6.
- Microsoft Corporation (2020), 'Microsoft Bug Bounty Program', https://www. microsoft.com/en-us/msrc/bounty?rtc=1.
- Mikolov, T., Chen, K., Corrado, G. & Dean, J. (2013), 'Efficient estimation of word representations in vector space'.
- Neuhaus, S., Zimmermann, T., Holler, C. & Zeller, A. (2007*a*), Predicting vulnerable software components, *in* 'Proceedings of the 14th ACM Conference on Computer and Communications Security', CCS '07, Association for Computing Machinery, New York, NY, USA, p. 529–540.

URL: https://doi.org/10.1145/1315245.1315311

Neuhaus, S., Zimmermann, T., Holler, C. & Zeller, A. (2007b), Predicting vulnerable software components, *in* 'Proceedings of the 14th ACM Conference on Computer and Communications Security', CCS '07, Association for Computing Machinery, New York, NY, USA, p. 529–540.
URL: https://doi.org/10.1145/1215245.1215211

URL: https://doi.org/10.1145/1315245.1315311

- Schuster, M. & Paliwal, K. K. (1997), 'Bidirectional recurrent neural networks', *IEEE Transactions on Signal Processing* **45**(11), 2673–2681.
- Shin, E. C. R., Song, D. & Moazzezi, R. (2015), Recognizing functions in binaries with neural networks, *in* 'Proceedings of the 24th USENIX Conference on Security Symposium', SEC'15, USENIX Association, USA, p. 611–626.
- Shin, Y., Meneely, A., Williams, L. & Osborne, J. A. (2011), 'Evaluating complexity, code churn, and developer activity metrics as indicators of software vulnerabilities', *IEEE Transactions on Software Engineering* 37(6), 772–787.
- Tien, C.-W., Chen, S.-W., Ban, T. & Kuo, S.-Y. (2020), 'Machine learning framework to analyze iot malware using elf and opcode features', *Digital Threats: Research and Practice* **1**, 1–19.
- Valeur, F., Mutz, D. & Vigna, G. (2005), A learning-based approach to the detection of sql attacks, *in* 'Detection of Intrusions and Malware, and Vulnerability Assessment: Second International Conference, DIMVA 2005, Vienna, Austria, July 7-8, 2005. Proceedings 2', Springer, pp. 123–140.

Wang, S., Liu, T. & Tan, L. (2016), Automatically learning semantic features for defect prediction, *in* 'Proceedings of the 38th International Conference on Software Engineering', ICSE '16, Association for Computing Machinery, New York, NY, USA, p. 297–308.

URL: https://doi.org/10.1145/2884781.2884804

- Wang, T., Wei, T., Gu, G. & Zou, W. (2010), TaintScope: A checksum-aware directed fuzzing tool for automatic software vulnerability detection, *in* '2010 IEEE Symposium on Security and Privacy', IEEE, pp. 497–512.
- White, M., Tufano, M., Vendome, C. & Poshyvanyk, D. (2016), Deep learning code fragments for code clone detection, *in* '2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)', IEEE/ACM, pp. 87–98.
- Wu, F., Wang, J., Liu, J. & Wang, W. (2017), Vulnerability detection with deep learning, *in* '2017 3rd IEEE International Conference on Computer and Communications (ICCC)', IEEE, pp. 1298–1302.
- Yamaguchi, F., Lottmann, M. & Rieck, K. (2012a), Generalized vulnerability extrapolation using abstract syntax trees, *in* 'Proceedings of the 28th Annual Computer Security Applications Conference', ACSAC '12, Association for Computing Machinery, New York, NY, USA, p. 359–368. URL: https://doi.org/10.1145/2420950.2421003
- Yamaguchi, F., Lottmann, M. & Rieck, K. (2012b), Generalized vulnerability extrapolation using abstract syntax trees, *in* 'Proceedings of the 28th Annual Computer Security Applications Conference', ACSAC '12, Association for Computing Machinery, New York, NY, USA, p. 359–368. URL: https://doi.org/10.1145/2420950.2421003
- Yamaguchi, F., Wressnegger, C., Gascon, H. & Rieck, K. (2013), Chucky: Exposing missing checks in source code for vulnerability discovery, *in* 'Proceedings of the 2013 ACM SIGSAC Conference on Computer & Communications Security', CCS '13, Association for Computing Machinery, New York, NY, USA, p. 499–510. URL: https://doi.org/10.1145/2508859.2516665
- Yang, X., Lo, D., Xia, X., Zhang, Y. & Sun, J. (2015), Deep learning for just-in-time defect prediction, *in* '2015 IEEE International Conference on Software Quality, Reliability and Security', IEEE, pp. 17–26.
- Yuan, D., Xie, Y., Panigrahy, R., Yang, J., Verbowski, C. & Kumar, A. (2011), Context-based online configuration-error detection, *in* 'Proceedings of the 2011 USENIX conference on USENIX annual technical conference', pp. 28–28.
- Zagane, M., Abdi, M. K. & Alenezi, M. (2020), 'Deep learning for software vulnerabilities detection using code metrics', *IEEE Access* **8**, 74562–74570.
- Zheng, Y. & Zhang, X. (2013), Path sensitive static analysis of web applications for remote code execution vulnerability detection, *in* '2013 35th International Conference on Software Engineering (ICSE)', IEEE, pp. 652–661.

Zhou, Y., Liu, S., Siow, J., Du, X. & Liu, Y. (2019), 'Devign: Effective vulnerability identification by learning comprehensive program semantics via graph neural networks', *Advances in neural information processing systems* **32**.