



An Empirical Evaluation of the Usefulness of Word Embedding Techniques in Deep Learning-Based Vulnerability Prediction

Ilias Kalouptsoglou^{1,2}(✉) , Miltiadis Siavvas¹ , Dionysios Kehagias¹ ,
Alexandros Chatzigeorgiou² , and Apostolos Ampatzoglou² 

¹ Centre for Research and Technology Hellas, Thessaloniki, Greece
{iliaskaloup,siavvasm,diok}@iti.gr

² University of Macedonia, Thessaloniki, Greece
achata.ampatzoglou@uom.edu.gr

Abstract. Software security is a critical consideration for software development companies that want to provide their customers with high-quality and dependable software. The automated detection of software vulnerabilities is a critical aspect in software security. Vulnerability prediction is a mechanism that enables the detection and mitigation of software vulnerabilities early enough in the development cycle. Recently the scientific community has dedicated a lot of effort on the design of Deep learning models based on text mining techniques. Initially, Bag-of-Words was the most promising method but recently more complex models have been proposed focusing on the sequences of instructions in the source code. Recent research endeavors have started utilizing word embedding vectors, which are widely used in text classification tasks like semantic analysis, for representing the words (i.e., code instructions) in vector format. These vectors could be trained either jointly with the other layers of the neural network, or they can be pre-trained using popular algorithms like word2vec and fast-text. In this paper, we empirically examine whether the utilization of word embedding vectors that are pre-trained separately from the vulnerability predictor could lead to more accurate vulnerability prediction models. For the purposes of the present study, a popular vulnerability dataset maintained by NIST was utilized. The results of the analysis suggest that pre-training the embedding vectors separately from the neural network leads to better vulnerability predictors with respect to their effectiveness and performance.

Keywords: Software security · Vulnerability prediction · Deep learning · Natural language processing · Word embedding vectors

1 Introduction

Vulnerability Prediction (VP) techniques aim to identify the software components that are more likely to contain vulnerabilities. Vulnerability prediction

models (VPMs) are typically built using machine learning (ML) techniques that use software attributes as input to differentiate between vulnerable and clean (or neutral) software components. Several VPMs have been proposed over the years, each of which uses different software factors as inputs to predict the presence of vulnerable components [1]. Text mining-based techniques have been found to be the most reliable, according to the bibliography [2] and have attracted the most of the recent research interest [3–10]. The first attempts in the field of vulnerability prediction using text mining, have focused on the concept of Bag-of-Words (BoW) as a method for predicting software vulnerabilities using the text terms and their respective appearance frequencies in the source code. Recently, researchers have shifted their focus from simple BoW to more complex approaches, investigating whether more complex textual patterns in the source code could lead to more accurate vulnerability prediction. In particular, the authors in [4, 5, 8] transformed the source code into sequences of word tokens and trained deep neural networks capable of learning sequences of data

When using sequences of tokens to identify software components with vulnerabilities, vulnerability prediction has a lot in common with text classification tasks such as sentiment analysis [11]. The word embedding vectors are commonly used in the field of text classification. Word embedding is a term used to describe the representation of words for text analysis, typically in the form of a real-valued vector that encodes the meaning of the word in such a way that words that are close in the vector space are expected to have similar meanings. Most of the studies about VP make use of the word embeddings [4–6, 12].

Tokens can be embedded into vectors in a variety of ways. One option is to use an embedding layer that is jointly trained with the vulnerability prediction task [13]. Another method is to use an external word embedding tool, such as word2vec [14], to generate vector representations of each token. One can also use the vectors that are already generated from these tools (e.g. word2vec, Glove [15], Fast-text [16]) based on natural language documents of billions of words. Finally, there is also the option to produce custom embedding representations.

The purpose of this study is to emerge the worth of the sophisticated embedding algorithms (e.g., word2vec, fast-text) in text mining-based vulnerability prediction showing their contribution to the effectiveness and the efficiency of the VPMs and to compare them with the use of a trainable embedding layer that updates its values during the training of the VP classifier. A dataset has been collected and an experimental analysis has been conducted by comparing the use of a simple embedding layer with the utilization of word2vec and fast-text algorithms. We also compare the word2vec and fast-text algorithms with each other. Finally, we compare our best model with a state-of-the-art model.

The remainder of the paper is organized as follows. Section 2 discusses related work regarding the utilization of word embeddings in the field of vulnerability prediction. Section 3 provides the theoretical background in order to familiarize the reader with the main concepts of the present work. Section 4 discusses the methodology that we followed, while Sect. 5 presents the results of our analysis. Finally, Sect. 6 wraps up the paper and discusses future research directions.

2 Related Work

Vulnerability prediction using text mining is very popular and has demonstrated promising results in the related literature [2, 4, 9, 10]. Initial research endeavors focused on the concept of BoW (i.e., occurrences of tokens) [2, 9]. Recent attempts focus on predicting the existence of vulnerabilities through learning more complex patterns from the source code. They consider the software components as sequences of tokens and train deep learning models capable of learning sequences, such as the Recurrent Neural Networks (RNNs) [4, 5]. The challenging part of these recent studies is to add syntactic and semantic meaning to the sequences of code tokens. Word embeddings are one of the most promising solutions.

The word embedding vectors have evolved into an integral part of the text classification tasks since Mikolov et al. [17] proposed two architectures for learning distributed representations of words. The authors in [18], conducted a comparative study between different ML algorithms including fast-text, Glove and word2vec, while in [19] a deep learning method is proposed utilizing the semantic knowledge provided by the word embeddings.

Word embeddings have already been used in the field of text mining-based vulnerability prediction. Dam et al. [8] mapped every code token with an index of their vocabulary and then they constructed an embedding matrix which contained a unique vector representation for every token of the vocabulary in the position that corresponds to the vocabulary index. In other words, the embedding matrix worked as a look-up table.

The authors in [5, 6] used the word2vec tool to generate embedding vectors for their vocabulary, while Zhou et al. [4] used the pre-trained word2vec vectors. Russel et al. in [12] created a vulnerability detection tool based on deep learning and capable of interpreting lexical source code. They conducted a comparative study between simple source code embedding using Bag-of-Words and more advanced code representations learned automatically by deep learning models inside the embedding layer. Fang et al. [20] proposed the fastEmbed model which is an extension of the fast-text algorithm. This way they developed a model for predicting the exploitability of software vulnerabilities on imbalanced datasets by understanding key features of vulnerability-related text.

To this end, it is quite clear that a lot of studies make use of word embedding vectors as a representative format for the source code's tokens (i.e., words). There are papers that refer the use of simple vector representations just in order to replace the text features [8], other papers that use the BoW methodology to represent the text in the source code [9], other studies that utilize the pre-trained embedding vectors produced by the pre-trained word2vec model [4], but most of them choose to encode the code tokens into embedding vectors trained on their own data [5, 6]. However, to the best of our knowledge, there is no study examining the difference between the internal embeddings that are trained in the embedding layer together with the classifier, and the external embeddings that are trained alone prior to the model's training. The former are part of the supervised learning of the model and update their weights through the Back-

propagation process [21], while the latter are trained once, using an advanced unsupervised algorithm, and then they can be saved for future use. Moreover, there is a need for an experimental analysis examining the improvement in terms of accuracy and performance that these word embeddings provide to the DL-based vulnerability predictors. In the present work, we attempt to address these open issues through an empirical analysis on a popular dataset. Furthermore, the present paper includes a comparison between two popular types of word embedding tools (i.e., word2vec, fast-text) as well as a comparison with a state-of-the-art BoW model.

3 Theoretical Background

3.1 Vulnerability Prediction Based on Text-Mining

Vulnerability Prediction purpose is to identify software hotspots that are more likely to contain software vulnerabilities. These hotspots are actually parts of the source code that require more attention by the software developers and engineers from a security viewpoint. When the VPMs are based on text-mining they are trained on datasets constructed by the words (i.e., tokens) that appear in the source code. BoW constitutes the simplest text-mining method. In BoW, the code is divided into text tokens, each one of which is accompanied by the number of its occurrences in the source code. So each word corresponds to a feature, and the frequency of that feature in a component adds up to the value of that feature for that component. Aside from BoW, text-mining includes the process of converting the source code into a list of token sequences for use as input to Deep Learning (DL) models capable of parsing sequential data (e.g., recurrent neural networks). The sequences of tokens constitute the input of the DL models that, during the training phase, try to capture the syntactic information included in the source code, and in the execution phase to predict the existence of vulnerabilities in the software components. Text-mining also uses Natural Language Processing (NLP) methodologies such as word2vec pre-trained embedding vectors to extract semantic information from tokens.

3.2 Word Embedding Vectors

Word embedding methods use a corpus of text to learn a real-valued vector representation for a predefined fixed-sized vocabulary [17]. The learning process is either collaborative with the neural network model on a task, or unsupervised, using document statistics. An embedding layer is a word embedding learned in conjunction with a neural network model on a specific natural language processing task, such as document classification. It necessitates cleaning and preparing the document text so that each word can be one-hot encoded. The model specifies the size of the vector space. The vectors are seeded with small random numbers. The embedding layer is used at the front end of a neural network and is fitted in a supervised manner using the Backpropagation algorithm. However, it can be

selected to be non-trainable. In this case, it has to be seeded with a pre-trained embedding matrix which has been trained using an external algorithm.

Mikolov et al. [17] proposed two model architectures for computing continuous vector representations of words. They showed that these representations were able to capture syntactic and semantic word similarities. Both architectures are neural network-based ones for learning the underlying word representations for every word. The first proposed model, called Continuous Bag-of-Words Model (CBOW), tends to find the probability of a word occurring in a context. Thus, it generalizes over all the different contexts in which a word can be used. The second architecture, called continuous skip-gram model, instead of predicting the current word based on context, attempts to maximize classification of a word based on another word in the same sentence. To be more specific, every current word is fed into a log-linear classifier with a continuous projection layer, which predicts words within a certain range before and after the current word.

Two of the most popular algorithms that can generate embedding vectors are the word2vec¹ and fast-text² models. Both of them are based on the two aforementioned architectures (i.e., CBOW, skip-gram). The difference between these tools lies in the fact that the word2vec considers each individual word to be the smallest unit for which a vector representation must be found, whereas fast-text considers a word to be formed by n-grams of character.

4 Methodology

4.1 Dataset

As part of the current work, we created several VPMs for two widely-used programming languages, C and C++ combined. We used a vulnerability dataset derived from two National Institute of Standards and Technology (NIST) data sources: the National Vulnerability Database (NVD³ and the Software Assurance Reference Dataset (SARD)⁴. This dataset contains 7651 class files, 3438 of which are classified as vulnerable and the remaining 4213 as clean. The dataset has been presented by Li et al. [5].

4.2 Pre-processing

Before the construction of vulnerability prediction models, appropriate pre-processing is required in order to bring the dataset in a form appropriate to be used by the investigated techniques. To this end, we gathered the source code files written in the C and C++ programming languages and used a variety of pre-processing techniques to convert the datasets into a series of words-tokens. All comments, as well as the header/import instructions that declare the use

¹ <https://radimrehurek.com/gensim/models/word2vec.html>.

² <https://radimrehurek.com/gensim/models/fasttext.html>.

³ <https://nvd.nist.gov/>.

⁴ <https://samate.nist.gov/SRD/index.php>.

of specific libraries in the class, were removed from the dataset. Subsequently, we removed the code-specific constants (i.e., numbers, literals, etc.), in order to make the produced sequences more generalizable. In particular, the numeric values (i.e., integers, floats, etc.) were then replaced by a unique identifier “numId\$”, while the string values and characters were replaced by a different unique identifier “strId\$”. All blank lines are also removed, and the text is finally transformed into a list of code tokens (i.e., new, char, strlen, etc.) in the order they appear in the source code. After data cleansing, these produced tokens are replaced by a unique integer (integer encoding process⁵) and these integers are mapped to one-hot vectors (one-hot encoding⁶).

4.3 Word Embedding Vectors Training

In this study, in order to embed the text representations to numerical vectors different from the one-hot vectors, the word2vec and fast-text tools were utilized. The two models provide both the CBOW and skip-gram architecture. We train these models with our dataset. Each software component constitutes a sequence of tokens and all the sequences of the dataset are used as the corpus for the training of the word2vec and fast-text models. These algorithms learn the syntactic and semantic relations between the code tokens and place them at the vector space. After training these embedding vectors for the words of the vocabulary then one can save them for future use, saving time of the training process. For the training of the embedding vectors, the parameters that were selected after tuning are listed in Table 1.

Table 1. The selected parameters for the training of word2vec and fast-text embedding vectors.

Parameters	Word2vec	Fast-text
Size	300	300
Window	40	40
min_count	1	1
Epochs	1	2

In Table 1, the parameter “size” is the dimension of the embedding vectors, the “window” refers to the maximum distance between a target word and words surrounding the target word while the term “min_count” refers to the minimum count of words to consider during training. The algorithm ignores the words with

⁵ <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.LabelEncoder.html>.

⁶ <https://scikit-learn.org/stable/modules/generated/sklearn.preprocessing.OneHotEncoder.html>.

4.4 Model Selection

In this analysis, various DL algorithms are used to create models that can distinguish between vulnerable and neutral source code files. As the input to the models consists of sequential data (i.e., series of tokens) we chose DL algorithms capable of handling sequences. The RNNs are the most suitable ones as language models [23]. Convolutional Neural Networks (CNNs) are used on code classification tasks, as well [4, 24]. Regarding the RNNs, there are several improved versions such as the Long-Short Term Memory networks (LSTMs) [25], the Gate Recurrent Units (GRUs) [26] and the Bidirectional LSTMs (BiLSTMs) [27], which can solve the vanishing gradient problem [28] that the original RNNs face. The hyper-parameters chosen for our RNN and CNN models are presented in Table 2. Their values were selected after consecutive tuning and re-evaluation.

Table 2. The selected hyper-parameters of the models.

Hyper-parameter name	Value for RNNs	Value for CNN
Number of layers	3 (Embedding-Recurrent-Dense)	3 (Embedding-Conv-Dense)
Number of hidden layers	1 (LSTM/GRU/BiLSTM)	1 (1D CNN)
Embedding size	300	300
Number of hidden units	300	128 filters
Kernel size	–	5
Weight initialization technique	Glorot Uniform (Xavier)	Glorot Uniform (Xavier)
Learning rate	0.01	0.01
Gradient descent optimizer	Adam	Adam
Batch size	64	64
Activation function	Relu	Relu
Output activation function	Sigmoid	Sigmoid
Loss function	Binnary Cross-entropy	Binnary Cross-entropy
Over-fitting prevention	Dropout = 0.3	–
Maximum epochs	100	100
Early stopping patience	10	10

4.5 Evaluation Metrics

Several evaluation metrics are available in the literature and are commonly used to assess the predictive effectiveness of the ML models. In the vulnerability prediction case, a special emphasis is placed on the Recall of the produced models,

because the higher the Recall of the model is, the more real vulnerabilities it predicts. Apart from the capability of the produced models to identify the great majority of vulnerable files contained in a software project, the volume of the produced False Positives (FP) (i.e., clean files marked as vulnerable) is important to consider because it is known to affect the models’ utilization in practice. The number of FP is closely related to the amount of manual effort required by developers to identify files that contain actual vulnerabilities. The lower the number of FP is, the higher the precision of the model. This fact emphasizes the significance of the f1-score, which represents the balance of precision and recall. However, because identifying vulnerable files at the expense of producing FP is more important in VP, we chose f2-score as our evaluation metric. The f2-score is a weighted average of precision and recall, with recall being more important than precision. It is equal to:

$$F_2 = \frac{5 \times \textit{precision} \times \textit{recall}}{4 \times \textit{precision} + \textit{recall}}$$

5 Results and Discussion

In this section, we present the results of our analysis and discuss the outcome of the experiments. Table 3 reports the evaluation results of the DL models that were built based on the sequences of tokens in the source code. This table sums up the results regarding the f2-score for all the RNN variations and CNN using word2vec or fast-text embeddings in contrast with the joint training of the embeddings with the neural network’s training. During the evaluation process, the ten-fold cross-validation process was employed eliminating the possibility of biased results.

Table 3. The f2-score of all the utilized methods after 10 fold cross-validation.

Model	S.E.L.	W2V 0	W2V 1	FastText 0	FastText 1
LSTM	77.98	85.11	88.38	84.92	88.66
BiLSTM	80.28	85.86	88.01	82.33	86.04
GRU	72.22	89.15	87.94	84.28	89.10
CNN	81.46	86.36	89.43	86.36	84.54
Average	77.99	86.62	88.44	84.47	87.09

S.E.L. = Simple Embedding Layer.

W2V: Word2vec.

0: CBOW.

1: Skip-gram.

From Table 3, we can see that the use of sophisticated word embeddings trained prior to the deep learning model is beneficial at each model case. The

average f2-score of the four models when using an algorithm for the generation of the word embedding vector is significantly bigger. Furthermore, it is clear that the skip-gram model is better than the CBOW in our dataset as it achieves greater average f2-score both at the word2vec and the fast-text case. Similarly, we notice that the word2vec method provides better f2-score, both at the CBOW and the skip-gram variation, compared with the fast-text embeddings. All the aforementioned findings lead to the conclusion that the skip-gram variation of the word2vec embeddings is the best choice for embedding the tokens of the source code of our dataset before giving them as input to the sequential deep learning model. An 11% increase in terms of f2-score when using word2vec embeddings compared with the trainable embedding layer is a significant improvement and indicates to the initial hypothesis that these sophisticated models are capable of capturing semantic and syntactic relationships between the words of the source code.

Furthermore, the training of the embedding vectors outside from the embedding layer (i.e., non-trainable embedding layer) is beneficial not only in accuracy but also in terms of performance. The training time of the DL models has decreased significantly. Table 4 sums up the results about the training time.

Table 4. The training time in milliseconds (ms) both in case of trainable embedding layer and in case of sophisticated embeddings trained independently of the neural network.

Model	S.E.L.	W2V 1
LSTM	13078	9090
BiLSTM	22596	18011
GRU	12025	8330
CNN	9774	4276

From Table 4, it is clear that the training times when having ready the embedding vectors are by far smaller compared with the case of joint training along with the rest layers. Another interesting note derived from Table 3 and 4 is the fact that the CNN model is more accurate than the RNNs and much faster as well.

Finally, another interesting question would be to examine whether the adoption of word embedding vectors lead to better vulnerability prediction models compared to the traditional (and simpler) BoW approach. For this purpose, we compare our best model that utilizes the word embedding concept to the best model that uses BoW and is trained and evaluated on the same dataset. In particular, in Table 5, we present the results of the comparison between the state-of-the-art BoW method, versus the CNN model with skip-gram word2vec vectors, which was found to be the best model in our previous analysis. In the case of BoW, we chose Random Forest (composed of 100 trees) as a classifier, based on bibliography [2,9]. From Table 5, it is observed that the f2-score is

greater in the case of using sophisticated embeddings. Actually, these word2vec vectors can be used only at token series models (i.e., CNN, RNN) and not in BoW, constituting a major drawback of the method.

Table 5. BoW versus CNN that uses the skip-gram word2vec representations.

Model	Accuracy	Precision	Recall	F2-score
BoW	88.69	90.40	85.80	86.66
CNN-W2V 1	88.25	86.21	90.31	89.43

6 Conclusion and Future Work

In this paper, we investigated the usefulness of the numerical representations of the source code words, with the aim of predicting vulnerabilities. We focused on examining whether the utilization of sophisticated (i.e., external) embedding vectors is beneficial in contrast with the training of the embedding vectors jointly with the vulnerability predictor. Moreover, a comparison between the CBOW and the continuous skip-gram architectures took place as well as a comparison between the word2vec and fast-text algorithms.

We showed that either the word2vec or fast-text methodologies provide better results than the trainable embedding layer which is trained along with the rest layers of the neural network. These vector representations seem able to capture semantic and syntactic relations between the words in the code and so they can be proved beneficial when training models on sequences of code tokens. The word2vec method proved to be superior to fast-text when applied in our dataset. Furthermore, the skip-gram model demonstrated better scores compared with the CBOW, both in cases of word2vec and fast-text. Another important advantage of these sophisticated vectors is the time reduction during the model training, as there is no need to train the embedding layer again. Last but not least, the CNN with trained word2vec embeddings, which appeared to be our best model, demonstrates higher f2-score than the BoW model.

There are several potential directions for future work. First of all, the present study was based on a dataset containing exclusively C/C++ code. We intend to replicate our study using software products written in other programming languages (e.g., Java, Python, etc.) to investigate the generalizability of the produced results. Furthermore, we aim to replicate our study by embedding the text features in a higher level of granularity (e.g., line or function level).

Acknowledgements. This work is partially funded by the European Union’s Horizon 2020 Research and Innovation Programme through IoTAC project under Grant Agreement No. 952684.

Appendix

In the Appendix Section we provide some extra Tables of the results produced by our study, including values for accuracy, precision and recall aside from the f2-score.

Simple Embedding Layer:

Model	Accuracy	Precision	Recall	F2-score
LSTM	75.61	74.69	79.60	77.98
BiLSTM	77.05	76.06	82.73	80.28
GRU	73.06	72.60	72.34	72.22
CNN	85.15	88.43	79.96	81.46

Word2vec Embeddings - CBOW:

Model	Accuracy	Precision	Recall	F2-score
LSTM	78.76	74.23	88.90	85.11
BiLSTM	80.62	75.79	88.99	58.86
GRU	84.55	79.52	92.02	89.15
CNN	86.60	86.16	86.52	86.39

Word2vec Embeddings - Skip-Gram:

Model	Accuracy	Precision	Recall	F2-score
LSTM	84.07	79.35	91.14	88.38
BiLSTM	84.51	80.26	90.31	88.01
GRU	83.44	78.51	90.74	87.94
CNN	88.25	86.21	90.31	89.43

Fast-Text Embeddings - CBOW:

Model	Accuracy	Precision	Recall	F2-score
LSTM	79.79	75.11	88.07	84.92
BiLSTM	76.20	71.15	85.79	82.33
GRU	77.58	72.71	88.53	84.28
CNN	86.08	84.98	86.82	86.36

Fast-Text Embeddings - Skip-Gram:

Model	Accuracy	Precision	Recall	F2-score
LSTM	83.54	78.23	91.79	88.66
BiLSTM	79.20	73.49	90.01	86.04
GRU	82.38	76.57	93.11	89.10
CNN	85.15	84.87	84.58	84.57

References

1. Siavvas, M., Gelenbe, E., Kehagias, D., Tzovaras, D.: Static analysis-based approaches for secure software development. In: Gelenbe, E., et al. (eds.) Euro-CYBERSEC 2018. CCIS, vol. 821, pp. 142–157. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-95189-8_13
2. Walden, J., Stuckman, J., Scandariato, R.: Predicting vulnerable components: software metrics vs text mining. In: 2014 IEEE 25th International Symposium on Software Reliability Engineering. IEEE (2014)
3. Chakraborty, S., Krishna, R., Ding, Y., Ray, B.: Deep learning based vulnerability detection: are we there yet. IEEE Trans. Softw. Eng. (2021)
4. Zhou, Y., Liu, S., Siow, J., Du, X., Liu, Y.: Devign: effective vulnerability identification by learning comprehensive program semantics via graph neural networks. arXiv preprint [arXiv:1909.03496](https://arxiv.org/abs/1909.03496) (2019)
5. Li, Z., et al.: Vuldeepecker: a deep learning-based system for vulnerability detection. arXiv preprint [arXiv:1801.01681](https://arxiv.org/abs/1801.01681) (2018)
6. Cao, S., Sun, X., Bo, L., Wei, Y., Li, B.: BGNN4VD: constructing bidirectional graph neural-network for vulnerability detection. Inf. Softw. Technol. **136**, 106576 (2021)
7. Pang, Y., Xue, X., Wang, H.: Predicting vulnerable software components through deep neural network. In: Proceedings of the 2017 International Conference on Deep Learning Technologies, pp. 6–10 (2017)

8. Dam, H.K., Tran, T., Pham, T.T.M., Ng, S.W., Grundy, J., Ghose, A.: Automatic feature learning for predicting vulnerable software components. *IEEE Trans. Softw. Eng.* **47**, 67–85 (2018)
9. Scandariato, R., Walden, J., Hovsepian, A., Joosen, W.: Predicting vulnerable software components via text mining. *IEEE Trans. Softw. Eng.* **40**, 993–1006 (2014)
10. Hovsepian, A., Scandariato, R., Joosen, W., Walden, J.: Software vulnerability prediction using text analysis techniques. In: *Proceedings of the 4th International Workshop on Security Measurements and Metrics*, pp. 7–10 (2012)
11. Medhat, W., Hassan, A., Korashy, H.: Sentiment analysis algorithms and applications: a survey. *Ain Shams Eng. J.* **5**(4), 1093–1113 (2014)
12. Russell, R., et al.: Automated vulnerability detection in source code using deep representation learning. In: *2018 17th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pp. 757–762. IEEE (2018)
13. Turian, J., Ratinov, L., Bengio, Y.: Word representations: a simple and general method for semi-supervised learning. In: *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pp. 384–394 (2010)
14. Rong, X.: word2vec parameter learning explained. *arXiv* (2014)
15. Pennington, J., Socher, R., Manning, C.D.: Glove: global vectors for word representation. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1532–1543 (2014)
16. Joulin, A., Grave, E., Bojanowski, P., Mikolov, T.: Bag of tricks for efficient text classification. *arXiv preprint arXiv:1607.01759* (2016)
17. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781* (2013)
18. Stein, R.A., Jaques, P.A., Valiati, J.F.: An analysis of hierarchical text classification using word embeddings. *Inf. Sci.* **471**, 216–232 (2019)
19. Ma, Y., Peng, H., Cambria, E.: Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In: *Thirty-Second AAAI Conference on Artificial Intelligence* (2018)
20. Fang, Y., Liu, Y., Huang, C., Liu, L.: Fastembed: Predicting vulnerability exploitation possibility based on ensemble machine learning algorithm. *Plos One* **15**, e0228439 (2020)
21. Goldberg, Y.: Neural network methods for natural language processing. *Synth. Lect. Hum. Lang. Technol.* **10**(1), 1–309 (2017)
22. Van der Maaten, L., Hinton, G.: Visualizing data using t-SNE. *J. Mach. Learn. Res.* **9**(11) (2008)
23. Sundermeyer, M., Oparin, I., Gauvain, J.L., Freiberger, B., Schlüter, R., Ney, H.: Comparison of feedforward and recurrent neural network language models. In: *2013 IEEE International Conference on Acoustics, Speech and Signal Processing* (2013)
24. Filus, K., Siavvas, M., Domańska, J., Gelenbe, E.: The random neural network as a bonding model for software vulnerability prediction. In: Calzarossa, M.C., Gelenbe, E., Grochla, K., Lent, R., Czachórski, T. (eds.) *MASCOTS 2020*. LNCS, vol. 12527, pp. 102–116. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-68110-4_7
25. Sundermeyer, M., Schlüter, R., Ney, H.: LSTM neural networks for language modeling. In: *Thirteenth Annual Conference of the ISCA* (2012)

26. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv (2014)
27. Schuster, M., Paliwal, K.K.: Bidirectional recurrent neural networks. *IEEE Trans. Signal Process.* **45**(11), 2673–2681 (1997)
28. Hochreiter, S.: The vanishing gradient problem during learning recurrent neural nets and problem solutions. *Int. J. Uncertainty* **6**, 107–116 (1998)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

