Predicting Organizational Cybersecurity Risk: A Deep Learning Approach

Benjamin Ampel
Department of Management Information Systems
University of Arizona
Tucson, Arizona, USA
BAmpel@email.arizona.edu

Abstract — Cyberattacks conducted by malicious hackers cause irreparable damage to organizations, governments, and individuals every year. Hackers use exploits found on hacker forums to carry out complex cyberattacks, making exploration of these forums vital. We propose a hacker forum entity recognition framework (HackER) to identify exploits and the entities that the exploits target. HackER then uses a bidirectional long short-term memory model (BiLSTM) to create a predictive model for what companies will be targeted by exploits. The results of the algorithm will be evaluated using a manually labeled gold-standard test dataset, using accuracy, precision, recall, and F1-score as metrics. We choose to compare our model against state of the art classical machine learning and deep learning benchmark models. Results show that our proposed HackER BiLSTM model outperforms all classical machine learning and deep learning models in F1-score (79.71%). These results are statistically significant at 0.05 or lower for all benchmarks except LSTM. The results of preliminary work suggest our model can help key cybersecurity stakeholders (e.g., analysts, researchers, educators) identify what type of business an exploit is targeting.

Keywords—Hacker forum, exploit, target identification, machine learning, deep learning

I. INTRODUCTION

Companies, governments, and individuals are increasingly being targeted by complex cyberattacks [1]. Cyberattacks are an attack launched by cybercriminals targeted against cyberinfrastructure [2]. As of 2018, each cyberattack can potentially lead to $7,010,000 in damages to an organization [3]. To help protect against such attacks, organizations are investing heavily in developing proactive Cyber Threat Intelligence (CTI) [4]. Automated data labeling techniques applied to hacker forums can provide proactive CTI, leading to better mitigation techniques for organizations [5]. Cyberattacks are generally performed by malicious hackers [2], who often congregate on forums to discuss tools (e.g., source code, binaries, executables) that can be used to perform devastating cyberattacks on governments, organizations, and individuals [6]. Traditional hacker forums contain tens of thousands of exploit source code that can be used for a potential cyber attack [7], with many of these tools also often mentioning the organization that the tool is targeting. An example is shown in Figure 1.
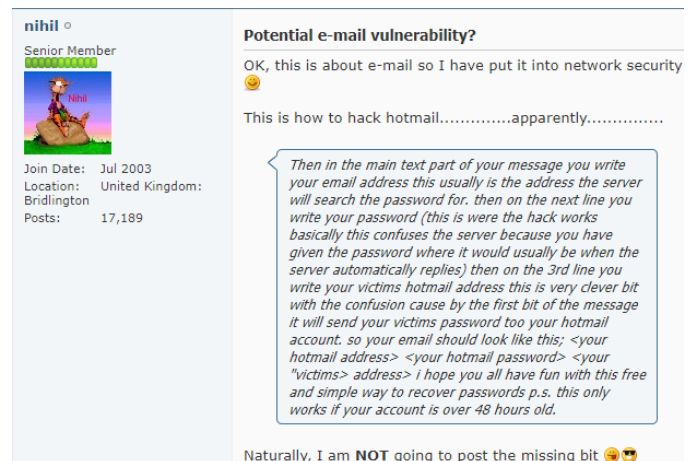

Fig. 1. Hacker Forum Post Mentioning an Organization (Hotmail)

In this study, we propose a novel entity resolution and classification framework, titled Hacker Entity Resolution (HackER). HackER automatically extracts organizational named entities (e.g., Hotmail) from hacker forum posts containing exploit source code to create a novel artifact that predicts what type of organization (e.g., open-source companies, video game companies) the exploit is targeting. Such an artifact would benefit cybersecurity stakeholders (e.g., analysts, researchers, educators) by providing more information on the types of companies that are targeted by hacker forum exploit source code.

The rest of the paper is organized as follows. First, we identify relevant literature related to hacker forum exploit analytics and bidirectional long short-term memory (BiLSTM) models. Second, we outline gaps found in the literature and we pose questions to address these gaps. Third, we present our novel research framework to answer the posed questions. Fourth, we discuss the practical implications of our work. Finally, we conclude this work and discuss future directions of the research.

II. LITERATURE REVIEW

For this study, we review two streams of literature: (1) hacker forum exploit analysis and (2) BiLSTM models. First, we research work done on hacker forum exploit analysis to study the methodologies used to identify and classify exploits
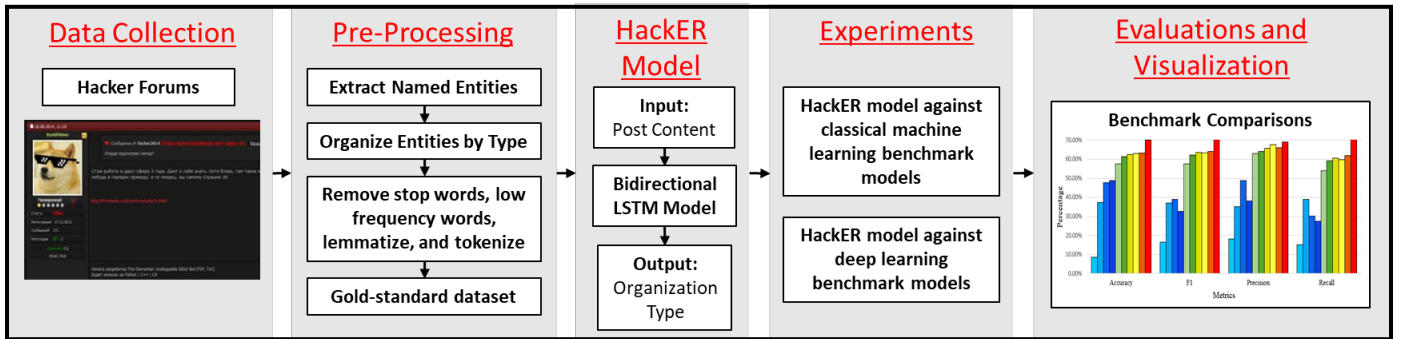
Fig. 2. Proposed HackER Framework

posted. Second, we review BiLSTM models to understand how to use the prevailing method for hacker forum exploit analysis for our target task of predicting organizational types targeted by posted exploits.

A. Hacker Forum Exploit Analysis

Hackers congregate at international hacker forums to share and discuss malicious tools that have been used in prior attacks against governments, companies, and individuals. Hacker forums generally contain millions of text-based posts. These posts are often noisy, making automated analysis methods non-trivial. Posts on hacker forums can have a significant effect on the occurrence of cyber-attacks [8], meaning there is great societal value in researching them.

Previous literature on hacker forums has focused on trend identification [9] and exploit categorization [5], [7], [10]–[12]. Within exploit categorization papers, we make three key observations. First, only two have used the exploit source code as an input feature [5], [7]. This is despite source code containing valuable information that can be extracted [13]. Second, while support vector machines (SVMs) were the prominent methodology used early on (2017-2018) [5], [11], [14], more recent literature (2018-2020) has adopted deep learning architectures for exploit identification [7], [10], [12]. The deep learning methods used are Recurrent Neural Network (RNN), Long-Short Term Memory (LSTM), and BiLSTM. As these models are prominent within hacker forum literature, their performance will be used as baselines in our proposed research.

B. BiLSTM models

As the BiLSTM algorithm has proven to be powerful in hacker forum exploit analysis [7] and hacker community analytics [15], we choose to review it for our task of predicting exploit organizational targets. BiLSTMs models can learn embeddings automatically from sequential text from both forward and backward contexts and preserve information from the future and the past [16]. Pre-trained word embeddings, like GloVe, can improve performance in classification tasks when used as the input to a BiLSTM layer [17]. GloVe is an unsupervised learning algorithm that obtains the vector representations of words based on their co-occurrence statistics [18].

III. RESEARCH GAPS AND QUESTIONS

Based on our review, we identified the following research gaps. The reviewed literature does an excellent job of extracting and categorizing the exploits that are being shared on hacker forums. However, there is a lack of work focusing on who (e.g., companies, individuals, governments, etc.) is being targeted by hacker forum exploits. Further, literature often discards source code when creating their analytic models. These gaps motivate the following research questions:

- How can named entity resolution be utilized to automatically and correctly extract organizational names from hacker forum exploit posts?
- How can we create a model that predicts the type of organization an exploit is targeting?

To answer these questions, we developed a novel research framework with five major components, shown in Figure 1.

IV. RESEARCH DESIGN

This study aims to build an automated alert system that informs organizations of exploits that are targeting them on international hacker forums. Our research component comprises of five main components (Figure 1): Data Collection, Pre-Processing, the HackER Model, Experiments, and Evaluations and Visualization.

A. Data Collection

Our collection contains posts from eleven prominent international hacker forums. Table 1 shows the seven features of metadata available from these hacker forums, including a description and examples.

TABLE I. DATA COLLECTION

Term	Description	Example
Title	Header for the exploit	inoERP 4.15 SQL Injection
Date	Date exploit was posted	26-Sep-19
Author	Person who posted the exploit	Alexandrovich Lyhin
Author Reputation	Respect of the author in the community	3/5 stars

Description	A short paragraph explaining what the code does	inoERP version 4.15 suffers from a remote SQL injection vulnerability.
Source Code	The code of the exploit itself	def generatePayload(query):#THIS FUNCTION IS INSECURE BY DESIGN b64_query = base64.b64encode(query); return os.popen("php -r \"echo base64_encode(serialize(base64_d ecode('" + b64_query + "')));\"").read()
Discussion	Comments that follow the posting of the code talking about the code	This still worked for me on Adobe 13.03

A crawler was developed for each forum, and the crawler was routed through the Tor network for anonymity. A depth-first search strategy was implemented for efficient parallel crawling through following different link stacks. This makes the process incremental, as a growing database of previously crawled links and dates is kept for each website, to ensure links are not visited or scraped twice. We summarize our research collection by name, forum language, number of posts, and number of source code snippets in Table 2.

TABLE II.   DATA COLLECTION

Name	Language	Posts	Source Code
0x00sec	English	9,161	397
Altenens	English	1,261,435	1,403
AntiChat	Russian	2,492,497	64,890
AntiOnline	English	291,914	2,063
Ciphers	English	51,612	2,207
Exelab	Russian	105,312	3,597
ExeTools	English	45,834	1,832
go4expert	English	62,103	5,800
KernelMode	English	29,755	934
WWHClub	Russian	1,492,156	53
WildersSecurity	English	2,571,053	2,096
Total	**2 Languages**	**8,412,832**	**85,272**

We chose to collect two different languages to ensure our data collection was representative of the international hacker community. In total, 85,272 exploit source code snippets were collected from 8,412,832 posts in the hacker forum testbed. Exploit source code was identified by using the DTL-EL tool for exploit labeling (Ampel et al., 2020). After collection, the data is stored in a MySQL database for further processing.

B. Pre-Processing

After collection, three steps are taken to pre-process the collected data. First, we utilize a pre-trained named entity recognition model from the Python package spaCy and apply it to hacker forum post content to extract organizational named entities. An example hacker forum post on an exploit source code snipper and the extracted organization are shown in Figure 2.

Fig. 3.   Example Hacker Forum Post with Highlighted (in blue) Organizational Entities

Second, we put each found organization into a bin (e.g., Mozilla into open source), and deleted bins with less than 100

organizational mentions. This is done to reduce our classification task to a reasonable number of labels. Third, source code was stripped of non-alphanumeric characters, lower-cased, lemmatized, tokenized, and padded to ensure proper lengths for all inputs. After the pre-processing step, a gold-standard testing dataset was constructed manually by removing irrelevant organizational mentions and keeping only five bins. Our gold-standard dataset is further detailed in Table 3.

TABLE III.   GOLD-STANDARD DATASET

Bin Label	Count	Percentage
Databases	1,780	34.17%
Software	1,351	25.94%
Open Source	961	18.45%
Mobile	673	12.92%
Video Games	445	8.52%
Total:	**5,210**	**100%**

The "Databases" bin will be the baseline of our model, as it is the majority label of our dataset (34.17% of the data). This bin contains companies such as MySQL, RethinkDB, and GenieDB. The software bin has many well-known companies such as Google and Oracle. Open source contains Mozilla and Mapbox. Mobile consists of RockYou, Showbucks, and Verizon. Finally, the video games bin contains companies such as Twitch, Oculus, and Zynga.

C. HackER Model

Consistent with hacker forum exploit analytics literature, we choose to use a BiLSTM model for our classification task. As shown in Figure 3, our proposed model consists of three layers, embedding, BiLSTM, and softmax.
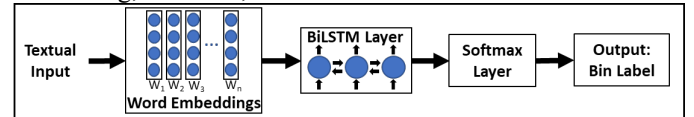

Fig. 4.   HackER Model Design

The input to our model is the tokenized exploit source code found in hacker forums. The word embedding layer uses the GloVe model [18] to capture the global and local statistics of our corpus to create word vectors. These word vectors are fed into the BiLSTM layer. A BiLSTM captures the past and future contexts of the input word vectors. This improves predictive performance over non-bidirectional approaches in long sequence-based text classification texts [17]. The past and future contexts are concatenated, and the final output is fed into the softmax layer. This layer produces a probabilistic score for each of our five labels and assigns the label with the highest score. This score is calculated by

$$\sigma(\vec{z})_i = \frac{e^{z_i}}{\sum_{j=1}^{K} e^{z_j}}$$

where $\vec{z}$ is the input vector given by the BiLSTM layer, i is the i-th bin class, and K is the number of total classes (in our case, five).

D. Experiments

Consistent with the computational design science paradigm [19], we evaluate our proposed artifact against benchmark models prominent in literature. Two experiments will be conducted to ensure the validity of our proposed approach. First, the proposed HackER will be evaluated against classical machine learning models (e.g., SVM, decision trees, naïve Bayes, logistic regression) found in hacker forum exploit analysis literature. Second, our model will also be benchmarked against prominent deep learning models (e.g., RNN, GRU, LSTM) found in related literature. All experiments use accuracy, precision, recall, and F1-score as metrics to evaluate. Each metric uses True Positives (TP), True Negatives (TN), False Positives (FP), and False Negatives (FN) to compute the metrics. The formulas for each metric are as follows:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}, Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}, F1 - score = \frac{2 \times Precision \times Recall}{Precision + Recall}$$

The uniformity in evaluation metrics allows us to compare the proposed HackER model against models commonly found in literature. All our models use the same 10-fold cross-validation for significance comparisons. Ten folds are consistent in literature [5] and is recommended by domain experts [20]. To evaluate statistically significant differences between HackER and benchmarks, we use one-tailed T-tests.

E. Evaluations and Visualizations

Consistent with the guidelines for deep learning-based information systems research [21], the experiments evaluated our proposed HackER (a BiLSTM model) against the proposed machine learning and deep learning benchmarks prominent in literature. The accuracy, F1-score, precision, and recall are summarized in Table 4 and visualized in Figure 4. Top scores are highlighted in bold-face.

TABLE IV.   EXPERIMENT RESULTS

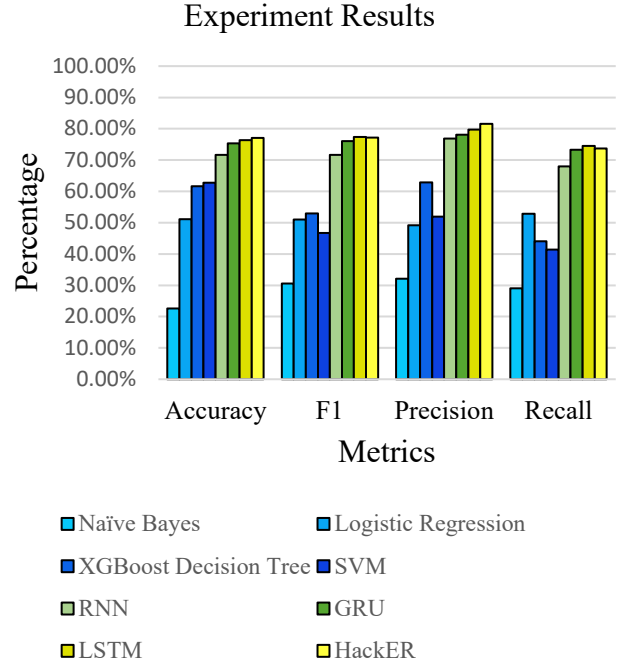| Experiment		Results			
Type	Model	Accuracy	F1-Score	Precision	Recall
Classical Machine Learning	Naïve Bayes	22.59% ***	29.08% ***	32.09% ***	30.58% ***
	Logistic Regression	51.16% ***	52.85% ***	49.13% ***	50.99% ***
	Decision Tree	61.65% ***	44.06% ***	62.87% ***	52.97% ***
	SVM	62.72% ***	61.38% ***	61.98% ***	66.68% ***
Deep Learning	RNN	71.64% ***	64.93% ***	76.89% ***	71.62% ***
	GRU	75.34% **	77.27% *	78.06% **	76.09% *
	LSTM	76.39% *	78.49%	79.77% *	**77.42%**
	HackER	**77.05%**	**79.71%**	**81.56%**	77.21%

Fig. 5.   Experiment Results

* indicates a statistically significant difference at $p < 0.05$, ** at $p < 0.01$, and *** at $p < 0.001$. We make two key observations about our experimental results. First, the DL model with the lowest F1-score (RNN) improved upon the best performing Classical ML model (SVM) by 3.55% (from 61.38% to 64.93%), and these results were significant at 0.001. Second, the F1-score of the HackER model (79.71%) outperformed all classical ML models (range 29.08% to 61.38%, significant at 0.001) and deep learning models (range 64.93% to 78.49%). The HackER improvement over the deep learning models was statistically significant at 0.05 or lower for all models except the LSTM. The LSTM also had an improvement over HackER in recall (from 77.21% to 77.42%), however this improvement was not statistically significant. These results suggest that using a BiLSTM for our HackER artifact

V. PRACTICAL IMPLICATIONS

The results of this work will have practical applications for key cybersecurity stakeholders (analysts and researchers). Analysts at various organizations can use the HackER model to gather proactive CTI on threats targeting their company or similar organizations. Researchers can build upon the framework or novel dataset to extract new and exciting information about hacker forum text.

VI. CONCLUSION AND FUTURE DIRECTIONS

In this study, we aimed to develop a novel artifact to automatically identify the type of organization an exploit posted on hacker forums targeted. Our results indicate that the HackER model can generalize well to the dataset and correctly label

exploit source code. The HackER model can be applied to exploit source code that does not have an organization mentioned with it to automatically extract the type of organization that is vulnerable to such an attack.

We have identified one promising direction for future work. In proactive CTI, it is vital to collect, identify, and mitigate potential cyber threats. Prior literature has made incredible strides in hacker forum exploit collection and identification. This work examines the targets of identified exploits. Future work can complete the proactive CTI methodology by creating a framework to automatically apply mitigation strategies to collected exploits. Such a framework could alert organizations about potential cyberthreats on hacker forums and inform them of mitigation strategies they can take before the exploit has the chance to be widely disseminated and used to cause damage. This direction can provide significant improvement to current proactive CTI efforts, ultimately contributed to a safer cyberspace for organizations, governments, and individuals.

VII. ACKNOWLEDGMENT

REFERENCES

[1] S. Samtani, H. Zhu, and H. Chen, "Proactively Identifying Emerging Hacker Threats from the Dark Web," *ACM Transactions on Privacy and Security*, vol. 23, no. 4, pp. 1–33, Aug. 2020, doi: 10.1145/3409289.

[2] V. Benjamin, J. S. Valacich, and H. Chen, "DICE-E: A Framework for Conducting Darknet Identification, Collection, Evaluation with Ethics," *MIS Quarterly*, vol. 43, no. 1, pp. 1–22, Jan. 2019, doi: 10.25300/MISQ/2019/13808.

[3] J. E. Lerums, L. D. Poe, and J. E. Dietz, "Simulation Modeling Cyber Threats, Risks, and Prevention Costs," in *2018 IEEE International Conference on Electro/Information Technology (EIT)*, May 2018, pp. 0096–0101, doi: 10.1109/EIT.2018.8500240.

[4] T. D. Wagner, K. Mahbub, E. Palomar, and A. E. Abdallah, "Cyber threat intelligence sharing: Survey and research directions," *Computers & Security*, vol. 87, Nov. 2019, doi: 10.1016/j.cose.2019.101589.

[5] S. Samtani, R. Chinn, H. Chen, and J. F. Nunamaker, "Exploring Emerging Hacker Assets and Key Hackers for Proactive Cyber Threat Intelligence," *Journal of Management Information Systems*, vol. 34, no. 4, pp. 1023–1053, Oct. 2017, doi: 10.1080/07421222.2017.1394049.

[6] V. Benjamin, W. Li, T. Holt, and H. Chen, "Exploring threats and vulnerabilities in hacker web: Forums, IRC and carding shops," in *2015 IEEE International Conference on Intelligence and Security Informatics (ISI)*, May 2015, pp. 85–90, doi: 10.1109/ISI.2015.7165944.

[7] B. M. Ampel, S. Samtani, H. Zhu, S. Ullman, and H. Chen, "Labeling Hacker Exploits for Proactive Cyber Threat Intelligence: A Deep Transfer Learning Approach," *2020 IEEE Conference on Intelligence and Security Informatics (ISI)*, no. November, 2020.

[8] W. T. Yue, Q.-H. Wang, and K.-L. Hui, "See No Evil, Hear No Evil? Dissecting the Impact of Online Hacker Forums," *MIS Quarterly*, vol. 43, no. 1, pp. 73–95, Jan. 2019, doi: 10.25300/MISQ/2019/13042.

[9] M. Schafer, M. Fuchs, M. Strohmeier, M. Engel, M. Liechti, and V. Lenders, "BlackWidow: Monitoring the Dark Web for Cyber Security Information," in *2019 11th International Conference on Cyber Conflict (CyCon)*, May 2019, pp. 1–21, doi: 10.23919/CYCON.2019.8756845.

[10] J. Grisham, S. Samtani, M. Patton, and H. Chen, "Identifying Mobile Malware and Key Threat Actors in Online Hacker Forums for Proactive Cyber Threat Intelligence," in *2017 IEEE International Conference on Intelligence and Security Informatics (ISI)*, Jul. 2017, pp. 13–18, doi: 10.1109/ISI.2017.8004867.

[11] I. Deliu, C. Leichter, and K. Franke, "Collecting Cyber Threat Intelligence from Hacker Forums via a Two-Stage, Hybrid Process using Support Vector Machines and Latent Dirichlet Allocation," in *2018 IEEE International Conference on Big Data (Big Data)*, Dec. 2018, pp. 5008–5013, doi: 10.1109/BigData.2018.8622469.

[12] R. Williams, S. Samtani, M. Patton, and H. Chen, "Incremental Hacker Forum Exploit Collection and Classification for Proactive Cyber Threat Intelligence: An Exploratory Study," in *2018 IEEE International Conference on Intelligence and Security Informatics (ISI)*, Nov. 2018, pp. 94–99, doi: 10.1109/ISI.2018.8587336.

[13] W. Wang *et al.*, "Reinforcement-Learning-Guided Source Code Summarization via Hierarchical Attention," *IEEE Transactions on Software Engineering*, pp. 1–1, 2020, doi: 10.1109/TSE.2020.2979701.

[14] I. Deliu, C. Leichter, and K. Franke, "Extracting cyber threat intelligence from hacker forums: Support vector machines versus convolutional neural networks," in *2017 IEEE International Conference on Big Data (Big Data)*, Dec. 2017, pp. 3648–3656, doi: 10.1109/BigData.2017.8258359.

[15] M. Ebrahimi, M. Surdeanu, S. Samtani, and H. Chen, "Detecting Cyber Threats in Non-English Dark Net Markets: A Cross-Lingual Transfer Learning Approach," in *2018 IEEE International Conference on Intelligence and Security Informatics (ISI)*, Nov. 2018, pp. 85–90, doi: 10.1109/ISI.2018.8587404.

[16] M. Thangaraj and M. Sivakami, "Text Classification Techniques: A Literature Review," *Interdisciplinary Journal of Information, Knowledge, and Management*, vol. 13, pp. 117–135, 2018, doi: 10.28945/4066.

[17] G. Liu and J. Guo, "Bidirectional LSTM with Attention Mechanism and Convolutional Layer for Text Classification," *Neurocomputing*, vol. 337, pp. 325–338, Apr. 2019, doi: 10.1016/j.neucom.2019.01.078.

[18] J. Pennington, R. Socher, and C. Manning, "Glove: Global Vectors for Word Representation," in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2014, pp. 1532–1543, doi: 10.3115/v1/D14-1162.

[19] J. F. Nunamaker, M. Chen, and T. D. M. Purdin, "Systems Development in Information Systems Research," *Journal of Management Information Systems*, vol. 7, no. 3, pp. 89–106, Dec. 1990, doi: 10.1080/07421222.1990.11517898.

[20] R. Kohavi, "A Study of Cross-Validation and Bootstrap for Accuracy Estimation and Model Selection," *Proceedings of the 14th international joint conference on Artificial intelligence*, pp. 1137–1143, 1995.

[21] S. Samtani, H. Zhu, B. Padmanabhan, Y. Chai, and H. Chen, "Deep Learning for Information Systems Research," no. Dl, pp. 1–55, 2020, [Online]. Available: http://arxiv.org/abs/2010.05774.