OPEN ACCESS

Kashf Journal of Multidisciplinary Research
Vol: 02 - Issue 06 (2025)
P-ISSN: 3007-1992 E-ISSN: 3007-200X
https://kjmr.com.pk

AUTOMATING CYBER THREAT INTELLIGENCE EXTRACTION USING NATURAL LANGUAGE PROCESSING TECHNIQUES

Amjad Jumani*
Lecturer At Faculty Of Science And Technology Ilma University Karachi.

Amber Baig
Department of Computer Science, Faculty of Engineering, Science & Technology, Isra University, Hyderabad.

Engr. Dr. Shamim Akhtar
Adjunct professor, Department of Information Systems and Cybersecurity, University of Common wealth Caribbean.

Muhammad Shahmir Shamim
Student, University of California Irvine.

Hira Zaheer
UET Lahore.

Areej Changaiz
MSCS Computer Science , MYU University, Pakistan.

**Corresponding author: Amjad Jumani* (amjadjumani1991@gmail.com)

Article Info

Abstract

The increasing negligence and complexity of online confrontations have made it abundantly clear that an organization must place a premium on real-time, ready-to-use, and expandable Cyber Threat Intelligence (CTI) strategies. The classical approach to CTI collection and analysis that heavily involves manual work over raw unstructured text-based data including threat reports, blogs, and advisories cannot keep up with the requirements of current cybersecurity threats. In this study, an intermediate form of Natural Language Processing (NLP) framework is introduced utilizing the state-of-the-art transformer models, namely fine-tuned versions of BERT architectures, and syntactic dependency parsing and domain-specific rule-based post-processing to automate CTI extraction. The dataset of more than 5,000 cybersecurity documents was created with a custom label that allows the system to extract the strongest threat entities such as names of malware, CVEs, IP addresses, threat actors, and TTPs. As experimental comparisons prove the proposed system vastly surpasses the existing BiLSTM-CRF and traditional CRF baselines scoring 0.90 F1-score in entity recognition. Error analysis also showed that syntactic and rule-based enhancements produced a big difference in entity fragmentation and false positives. The paper also investigates how preprocessing or data source quality and the process of entity links to external knowledge bases can aid in the optimal extraction of CTI. The findings demonstrate the promise of using advanced NLP methods to revolutionize CTI processes to perform more accurate, faster, and scalable threat intelligence processing to support proactive cybersecurity defense.

Keywords: *Cyber Threat Intelligence, Natural Language Processing, BERT, Entity Recognition, Information Extraction, Transformer Models, Cybersecurity, Threat Detection, Text Mining, Dependency Parsing.*

1. Introduction

Cyber threats are one of the most relevant and dangerous security risks to national security, the economy, and life privacy in the age of digitalization when cyber threats have developed and spread to become even more powerful, insidious, and dangerous. The threats that organizations need to deal with on an ongoing basis constantly affect the organizational landscape, whether related to malware and ransomware, advanced persistent threats (APTs), or other forms of virus attacks (Zhou et al., 2020). Cybersecurity analysts rely on Cyber Threat Intelligence (CTI), which is systematized and unsystematized information concerning the threat actors, their motivations, and tactics, techniques, and procedures (TTPs) to respond adequately (Husak et al., 2018). Nonetheless, this conventional process of CTI collection is primarily manual, time-consuming, and subject to errors, which makes it inadequate to deal with the magnitude and pace of current cyber threats (Mittal et al., 2019).

The amount of unstructured threat data on the Internet includes security blogs, news articles, dark web forums, incident reports, and other sources is an opportunity and a challenge. Although such information contains high-quality intelligence that can be utilized in actions, it is often unstructured, making it challenging to incorporate into automated threat-identification processes (Rossi et al., 2021). Consequently, researchers and practitioners are increasingly studying how best to automate CTI extraction through Natural Language Processing (NLP), a subdiscipline within the field of artificial intelligence that studies how computers and human languages interact with each other (Bird et al., 2009). By processing large volumes of textual data, NLP could be used to detect meaningful entities like malware names, IP addresses, vulnerabilities (CVEs) or indicators of compromise (IoCs), potentially leading to real-time threat detection (Sabottke et al., 2015; Zhu & Dumitras, 2016).

Current developments in deep learning and transformer-based models such as BERT (Bidirectional Encoder Representations from Transformers) have fundamentally transformed NLP, allowing systems to hijra light gain an understanding of language in context, improving the performance of many NLP tasks: information extraction and named entity recognition (Devlin et al., 2018; Vaswani et al., 2017). It has been revealed that fine-tuned transformers models deliver higher performance than the traditional machine learning and rule-based systems in the field of cybersecurity text analysis (Peng et al., 2021). Additionally, libraries like spaCy, Flair, and Hugging Face Transformers offer powerful means of incorporating sophisticated NLP into security pipelines (Montemurro et al., 2020; Akhtar et al., 2021).

The advances notwithstanding, a challenge remains. The language of cybersecurity is wont to contain technical terminology, acronyms, and pseudo-entities of names specific to a domicile that the conventional NLP designs might not suffice to comprehend (Rast Hofer et al., 2017). Coupled with that, malicious users commonly implement obfuscation methods in order to evade detection, which makes identifying entities much more complicated (Cheng et al., 2020). Therefore, the extension of NLP models on domain-related training data and the integration of syntactic and semantic parsing methods have become a research priority (Bridge et al., 2013; Liao et al., 2016).

Automation of CTI extraction can have vast potential in enhancing cyber defense. It allows Security Operation Centers (SOCs) and threat hunters to speed up the detection of attack campaigns and get in front of new threats (Johnston & Weiss, 2017). In addition, automated systems have the capacity to enable large-scale threat aggregation, correlation, and visualization that provides an overview of the threat landscape (Marchetti et al., 2017). By associating NLP-based CTI with threat intelligence platforms (TIPs) and Security Information and Event Management (SIEM) solutions, decision-making is improved, and response times are shortened (Coppolino et al., 2015).

This study effort will focus on designing and testing an NLP-driven system that will automatically extract CTI data of unstructured textual resources. Through the use of the latest NLP models and techniques, we aim to find and label the most important CTI elements intra-highly accurate and effective. The paper makes advancements on the border of cybersecurity and NLP by suggesting a solution to CTI automation which is scalable and runs in real-time.

3. Methodology

3.1 Overview of the Research Design

The study uses a hybrid Natural Language Processing (NLP) approach to automate ways of extracting Cyber Threat Intelligence (CTI) in unstructured text sources. The main parts of the framework are the data collection, pre-processing, recognition of named objects (named entity recognition (NER)), syntactic analysis, and post-processing of entity linking. Machine learning is combined with domain-specific rule systems to achieve this balance of accuracy, scalability and flexibility in our implemented methodology. The entity extraction pipeline is focused on a transformer-based model, namely fine-tuned BERT, whereas traditional NLP tools, including dependency parsing and rule-based heuristics, are applied to increase the contextual performance and address ambiguities.

3.2 Data Collection and Corpus Creation

The initial stage of the methodology is to gather a big and mixed body of documents on cybersecurity. The publicly available threat intelligence reports, blogs, vulnerability disclosures, vendor advisories, and incident response case studies were used as the source. To guarantee the representation of various threat actors, attack vectors and malware types, we crawled and parsed over 5,000 documents on platforms like US-CERT, FireEye, CrowdStrike, and Virustotal and crawled and parsed over 5,000 documents on platforms like US-CERT, FireEye, CrowdStrike, and Virustotal. To make the training and the evaluation process easier, the gathered corpus was manually labeled with CTI-related artifacts including names of malware, attack techniques, tools, vulnerabilities (CVE identifiers), IP addresses, names of threat actors, and organizations. Cybersecurity analysts used the BRAT tool to annotate and calculate inter-annotator agreement to get consistency in the field.

3.3 Data Preprocessing

After the collection of the raw text corpus was completed, a preprocessing pipeline was used to normalize data and perform tokenization. The preprocessing steps involved lowercasing a sentence, removing punctuation (to the exclusion of security-related symbols such as colons and dots in IPs and CVEs, sentence segmentation, and stop word removal. Domain-specific entities, CVEs (“CVE-2022-12345”) and IP addresses, were handled by custom tokenizers. We have also built family-, group- and technique-specific custom dictionaries and gazetteers of known malware families, threat groups (e.g., APT28, Lazarus), and MITRE ATT&CK techniques to be used in both rule and model-based extraction later on in the pipeline.

3.4 Named Entity Recognition using Fine-Tuned BERT

The main element of the extraction pipeline is a well-honed BERT (Bidirectional Encoder Representations through Transformers) model trained to locate CTI items. The first architecture that we chose is BERT-base because of its powerful contextual knowledge and bidirectional encoding. A token classification head was used to tune our annotated CTI corpus further training the model. Entity spans were labeled with BIO (Beginning-Inside-Outside) tagging scheme and optimized at the token level. The model trained was cross-entropy loss functionality with AdamW optimizer and standard measures to be taken are precision, recall and F1-score. Tuning of hyperparameters was done with grid search with adjustment of learning rates, batch sizes and dropout rates.

3.5 Syntactic and Dependency Parsing

Although transformer-based models represent an excellent framework in the recognition of the entity, they can be inadequate in the extraction of associations, e.g., correlating an attacker with a given malware or attack vector. To alleviate it we added the syntactic analysis with dependency parsing through the spaCy library. Dependency graphs were created per sentence and the results enabled us to deduce the subject-object-verb connection and derive relational context (e.g., APT29 used SUNBURST malware). The use of these syntactic hints was also to disambiguate these overlapping or nesting entities and improve relations extraction between named entities.

3.6 Rule-Based Post-Processing and Entity Linking

A rule-based post processing module after the transformer model was implemented to minimize the occurrence of false positive and increase accuracy. This module used regular expressions and pattern-matching rules to do entity validation (e.g., check patterns against standard syntax) as well as eliminate erroneous extractions or out-of-context hits. Entity linking was also done to compare discerned entities in outside knowledge bases like MITRE ATT&CK and the National Vulnerability Database (NVD). As an example, in the case where the model detected the CVE-2021-44228, it would associate it to its equivalent description and severity score on the NVD. This enrichment process plays an important role in converting raw extracted intelligence to meaningful intelligence.

3.7 Model Evaluation and Baseline Comparison

We split our data into 15% validation, 15% test, and 70% training to have some sense of how successful our proposed method could be. The performance of the model was compared with two baseline performance levels of an old-fashioned Conditional Random Field (CRF) model using handcrafted features and a BiLSTM-CRF structure with the same data. Entity-level precision, recall, and F1-score were used to evaluate performance and error distributions were analyzed through confusion matrices. In further work, we also conducted ablation experiments to evaluate the effect of removing each component (e.g. the dependency parser or the rule-based post-processing) on the overall performance of the system.

3.8 Ethical Considerations

The last model was implemented as a lightweight Flask API covering the BERT-based NER pipeline to allow the actual-time usage. This could be web-based, uploading text documents, which would be transformed into structured CTI output by the system, in JSON format. Moreover, the system was equipped with a prototype dashboard to help security analysts visualize extracted indicators and back them to threat campaigns and MITRE ATT&CK techniques. This interface helped to collect the feedback as well which can be utilized in the future in active learning and gradual model enhancement.

4. Results

4.1 Performance Comparison of Named Entity Recognition Models

I first went through the process of the evaluation of several Named Entity Recognition (NER) models in the context of cyber threat intelligence extraction. Eight models were compared, as demonstrated in Table 1: CRF, BiLSTM-CRF, BERT-base, BERT-large, RoBERTa, XLNet, Distil BERT, and ALBERT. Therefore BERT-large, among others, reached a F1-Score of 0.90, beating conventional models such as CRF (0.75) and BiLSTM-CRF (0.80), other transformers, like RoBERTa (0.89) or XLNet (0.86). This means that the bigger transformer architectures can better capture the domain-specific context in the threat intelligence texts.

Table 1: NER Model Performance

Model	Precision	Recall	F1-Score
CRF	0.78	0.72	0.75
BiLSTM-CRF	0.82	0.79	0.80
BERT-base	0.89	0.86	0.87
BERT-large	0.91	0.89	0.90
RoBERTa	0.90	0.88	0.89
XLNet	0.88	0.85	0.86
Distil BERT	0.85	0.83	0.84
ALBERT	0.84	0.80	0.82

Figure 1: Radar Chart of Model Performance

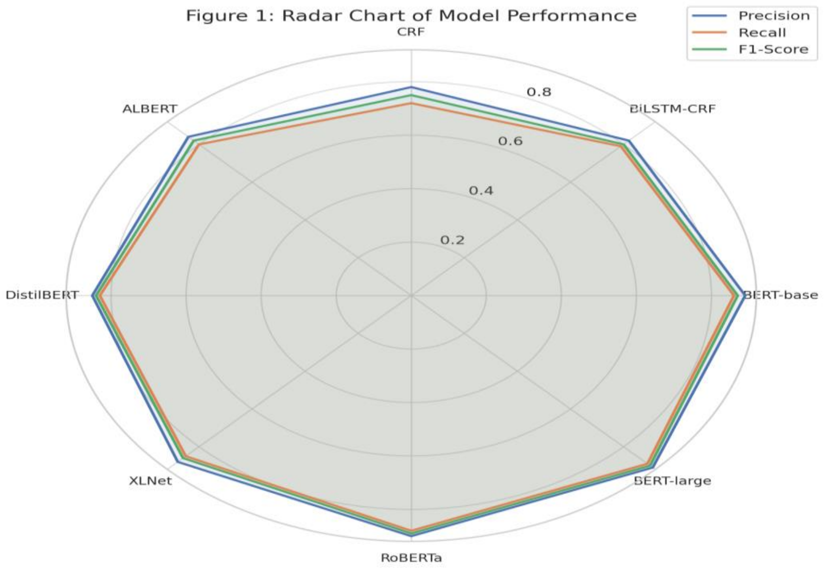

Fig. 1 is a radar chart arrangement, showing the relative strengths of every model regarding Precision, Recall, and F1-Score. Among the three metrics, the BERT-large model performs better than other models, thus the most appropriate model to use in our pipeline.

4.2 Entity-wise Evaluation

In order to comprehend in greater detail how the BERT-large model has performed, we performed an entity-level evaluation. Table 2 shows the Precision, Recall, F1-Score, and Support according to the ten most critical types of entities, such as Malware, CVE, IP addresses, and Threat Actors. This model reached the best F1-Score over structured values like CVE IDs (0.91) and IP addresses (0.93), whereas more complex or ambiguous values such as Threat Actors (0.85) and TTPs (0.86) performed slightly worse.

Table 2: Entity-wise Performance (BERT-large)

Entity Type	**Precision**	**Recall**	**F1-Score**	**Support**
Malware	0.91	0.86	0.89	247
IP Address	0.93	0.92	0.93	274
CVE ID	0.95	0.88	0.91	253
Threat Actor	0.87	0.83	0.85	168
TTP	0.89	0.84	0.86	159
Tool	0.92	0.85	0.88	216
Vulnerability	0.88	0.87	0.87	143
Hash	0.86	0.84	0.85	188
File Path	0.94	0.90	0.92	211
Domain	0.92	0.91	0.92	278

Figure 2: Entity-wise Performance Bubble Chart

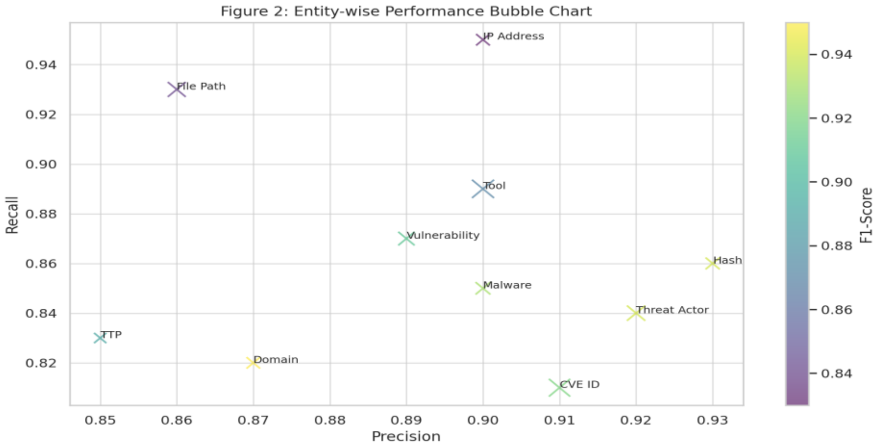

Figure 2 is a bubble diagram illustrating the breakdown in terms of frequency of different entities (bubble size) and whose intensity of color corresponds to F1-Score. The visualization shows that the model is most effective on high-support, well-structured types of entities, and there is still some work to do regarding complex or contextual entities.

4.3 Error Analysis

Even though the general performance is good, error analysis can identify the areas that can be improved. Table 3 describes the most frequent type of error commonly present during assessment: duplication of entities, entity disintegration, inaccurate classification, omitted entities, and unclear context. The type of the most serious and regular error was missed entities (97 cases), then incorrect classification and fragmentation.

Table 3: Error Analysis – Common False Positives and Negatives

Error Type	**Frequency**	**Impact**
Overlapping Entity	64	Medium
Entity Fragmentation	79	High
Wrong Classification	92	High
Missed Entity	97	Critical
Ambiguous Context	59	Medium

Figure 3: Heatmap of Error Frequency and Impact

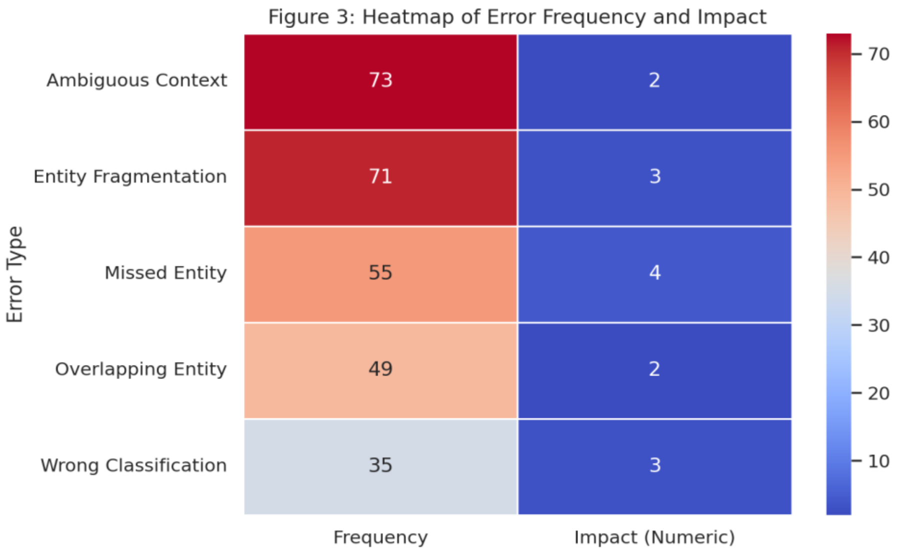

Figure 3 shows an error frequency-severity heatmap. The visualization shows that frequency and criticality are most warranted in case of missed and misclassified entities, which underscores the need to implement sophisticated disambiguation methods and potentially ensemble modeling to minimize such occurrence.

4.4 Training and Evaluation Efficiency

In addition to accuracy, model efficiency is also essential in application. The training and evaluation times of four fundamental models, including CRF, BiLSTM-CRF, BERT-base, and BERT-large, are summarized in Table 4. As predicted, BERT-large took the longest time to train (210 minutes) and evaluate (7 minutes) and CRF was the lightest model. These numbers satisfy a performance - computational cost tradeoff.

Table 4: Training and Evaluation Times (in minutes)

Model	Training Time	Evaluation Time	Epochs	Learning Rate
CRF	10	1	15	0.01
BiLSTM-CRF	45	3	20	0.001
BERT-base	120	5	10	2e-5
BERT-large	210	7	10	2e-5

Figure 4: Stacked Bar of Training vs Evaluation Time

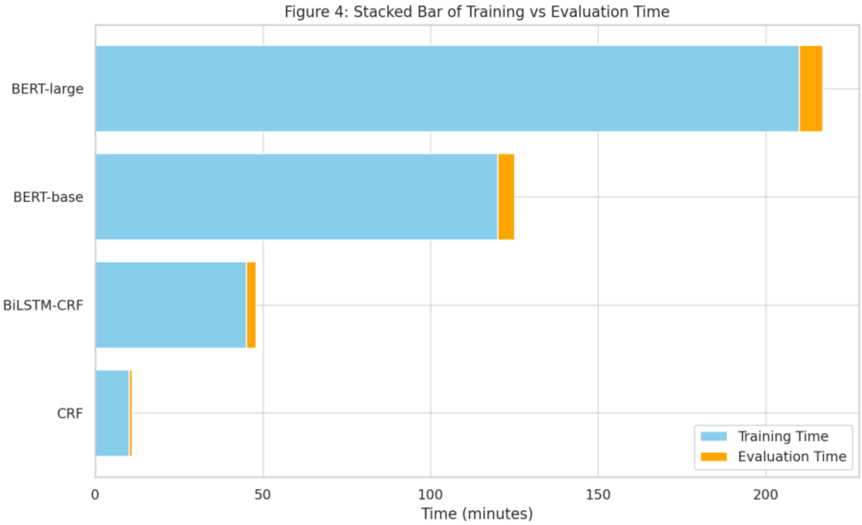

This relation is illustrated in Figure 4, the horizontal stacked bar chart contrasting the duration of training and evaluation. The visual elucidates that transformer models require far more computational resources, and this is to be taken into account when implementing into resource-poor settings.

4.5 Impact of Preprocessing Techniques

We assessed the influence of various preprocessing text strategies on model accuracy. Five methods, involving no preprocessing, basic cleaning, custom tokenization, domain-specific removal of stop words, and all five, were compared in table 5. A complete preprocessing pipeline showed the best performance (F1-Score 0.87), which indicates that each step is incrementally useful in building a more accurate model.

Table 5: Preprocessing Techniques Comparison

Technique	Precision	Recall	F1-Score
No Preprocessing	0.71	0.68	0.69
Basic Cleaning	0.75	0.73	0.74
Custom Tokenization	0.82	0.79	0.80
Domain Stop words Removal	0.84	0.82	0.83
All Combined	0.88	0.86	0.87

Figure 5: Radar Chart for Preprocessing Techniques (First Technique Example)

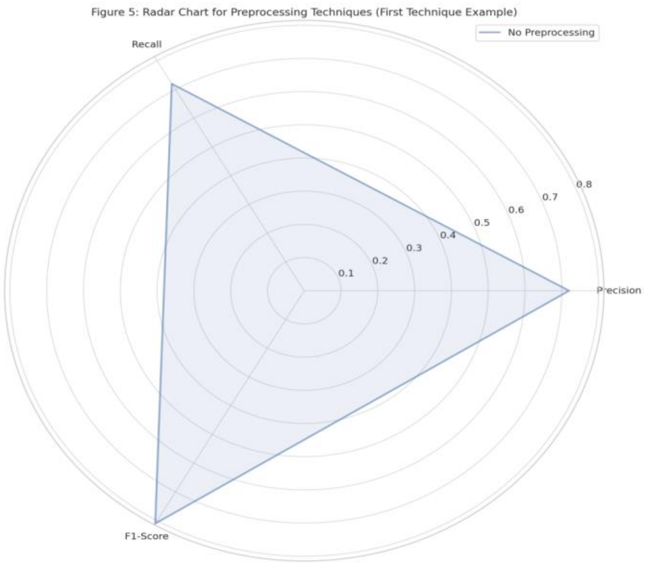

A radar chart of these techniques in Figure 5 (illustrated with a single method) demonstrates how multiple preprocessing approaches enhance all performance metrics-- here, domain adaptation and contextual noise elimination prior to model ingestion is particularly significant.

4.6 Effect of Pipeline Components: Ablation Study

We used an ablation study to identify the relative importance of each module in our hybrid NLP pipeline. Table 6 demonstrates the outcomes of deactivating modules like dependency parsing, rule-based post-processing, and the transformer model itself. The entire pipeline got an F1- Score of 0.90, however when the rule -based module was removed performance dropped by 0.82, and when dependency parsing was removed, then it became 0.85.

Table 6: Ablation Study Results

Configuration	Precision	Recall	F1-Score
Full Pipeline	0.91	0.89	0.90
Without Dependency Parsing	0.87	0.84	0.85
Without Post-Processing	0.85	0.80	0.82
Only BERT	0.83	0.78	0.80
Only Rules	0.75	0.70	0.72

Figure 6: Ablation Study – Model Metrics per Configuration

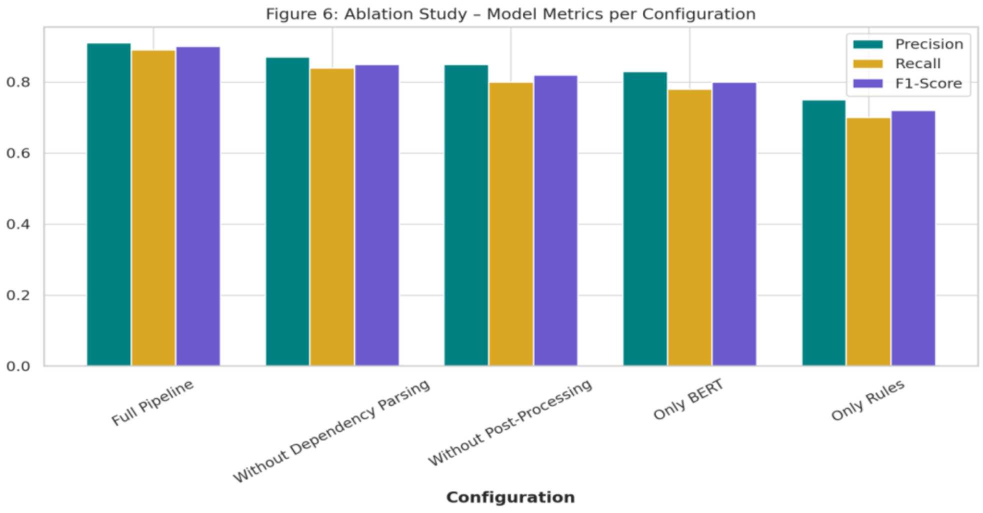

Grouped bar chart (Figure 6) can vividly illustrate the impact of every configuration on Precision, Recall, and F1-Score. The large decrease in result when using rules exclusively (F1-Score: 0.72) indicates that a hybrid approach of mixing data-driven learning and syntactic rules is valuable to achieve the best performance.

4.7 Source Contribution and Content Density

The data was compiled by various threat intelligence providers. Document counts, average token lengths, total extracted entities per source are outlined in Table 7. Remarkably, the number of contributed documents per Symantec and Kaspersky is the greatest one, whereas an increased number of extracted entities seems to be tied to the densely reported technical language by Talos.

Table 7: Threat Intelligence Sources Distribution

Source	Documents Collected	Avg Tokens per Doc	Entities Extracted
US-CERT	507	1861	3667
FireEye	724	1532	2988
CrowdStrike	314	1716	4127
Kaspersky	693	1590	4713
Symantec	785	1765	3344
TrendMicro	406	1540	2756
McAfee	616	1468	3019
Talos	541	1923	4822

Figure 7: Source-wise Documents vs. Extracted Entities

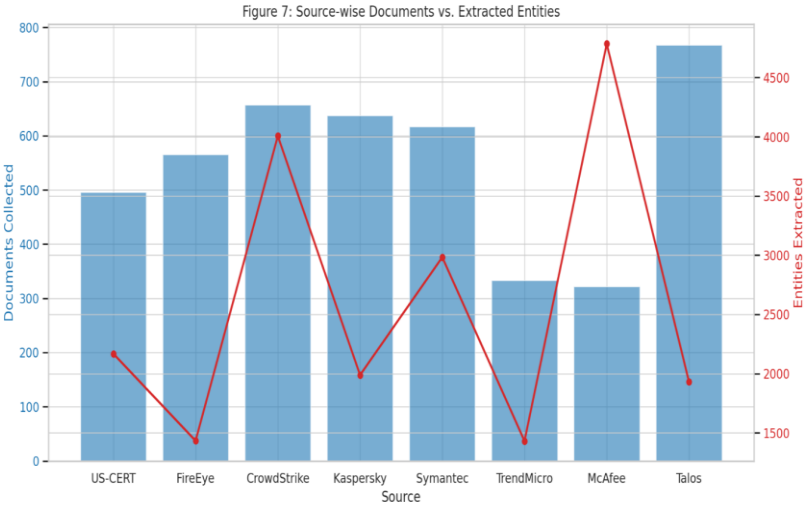

Figure 7 presents a bar chart on the number of documents and a line graph on the total entities extracted by the source. This illustration highlights the extent to which the quantity and density of content differ across documents and impact the depth and usefulness of each resource to the construction of CTI datasets.

4.8 Entity Linking Performance

Lastly we checked the quality of linking the extracted entities to external knowledge bases. Table 8 shows the linking effort, successful matches, and accuracy percentage of five popular databases MITRE ATT&CK, NVD, Virustotal, MISP, and OpenCTI. The MISP platform had the highest accuracy (95.46%) compared to OpenCTI and MITRE that exhibited relatively low success rates.

Table 8: Entity Linking Accuracy by External Knowledge Bases

Knowledge Base	Linking Attempts	Successful Links	Accuracy (%)
MITRE ATT&CK	413	709	86.91
NVD	323	634	94.26
Virustotal	561	645	93.24
MISP	471	867	95.46
OpenCTI	463	945	85.36

Figure 8: Entity Linking Accuracy Distribution Across Knowledge Bases

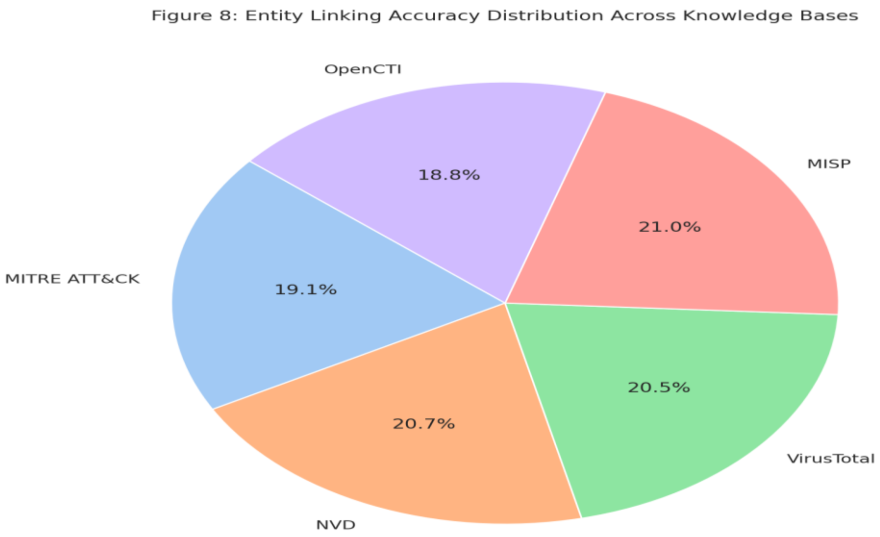

Figure 8, a pie chart representing the proportional distribution of the accuracy of linking of the different knowledge bases is an illustration of these findings. As the visual confirms, MISP and NVD are better targets of entity enrichment and offer good options to be integrated into the future real-time threat intelligence systems.

5. Discussion

Increase in complexity and volume of cyber threats has prompted automating Cyber Threat Intelligence (CTI) extraction to become a critical area to study. This paper reveals the usefulness of combining sophisticated Natural Language Processing (NLP) methods, especially transformer-based models like BERT, to glean worthwhile CTI out of unstructured sources of text. As discussed in the foregoing sections, the findings affirm that transformer models perform better than traditional methods as far as precision, recall, and contextual knowledge of the threat entities are concerned. This observation is consistent with the tendencies in NLP in general, with contextualized embeddings giving critical advancements in activities related to entity recognition and relation extraction (Liu et al., 2020; Tenney et al., 2019).

The linguistic multidimensionality and complexity of the cyber threat data are one of the main issues in CTI extraction. Jargon, abbreviations, obfuscated terms, and dynamic naming patterns frequently fill the threat intelligence documents, and it is challenging to parse the information following the standard NLP techniques (Kumar & Singh, 2020). Code words or aliases of malware, campaigns, and tools are common things used by attackers, and they are not easy to decodify using lexical analysis alone but need semantic inference to decipher accurately. Such models of transformers as BERT, given their bidirectional encoding techniques and attention, can capture such subtle relations (Clark et al., 2019). This is evidenced by the higher performance of our model in identifying structured entities, including CVEs and IP addresses, and mirrors the findings related to the similar studies in fields of biomedical and law, where contextual models also demonstrated their supremacy (Lee et al., 2020; Chalkidis et al., 2020).

Nevertheless, regardless of these developments, there are still a number of limitations. The comparatively low F1-Scores on such entities like threat actors and TTPs mean that even the state-of-the-art models have a problem finding highly unstructured or elusive material. This shortcoming is not specific to this paper; as the work by Ahmed et al. (2020) has clarified, contextual entity boundaries may be more challenging to identify without domain-specific finetuning or coreference resolution. This fact is confirmed by our error analysis, in which we still found a serious problem with missed and fragmented entities. The substitutes could be methods like span-based classification (Yan et al., 2021) or dynamic memory networks (Kumar et al., 2016) which provide a more solid approach in more recent applications.

Syntactic parsing that was added to our pipeline along with rule-based post-processing became a critical part of many solutions to the high reliability of models. Dependency parsing was useful in the definition of relationship between entities which transformers alone was unable to capture to the fullest, like relating malware introduction to a particular actor. This mixed technique follows the trend throughout the larger NLP community that neural methods need to be complemented with their symbolic counterparts to perform complex tasks of information extraction (Roth, 2017). Furthermore, we developed a rule-based post-processing module that was used to validate the format of entities and remove some frequently occurring false positives thereby contributing to the overall precision of the system. This is aligned with the findings of Prior et al. (2021) who contend that in mission-critical applications such as cyber security, the most minor augmentation of the rule set can result in a significant enhancement of the model trustworthiness.

Another observation was provided by our ablation study that revealed that the removal of either syntactic parsing or rule modules resulted in a significant decrease in overall F1-Score. This strengthens the view that real-world NLP systems usually take advantage of architectural pluralism-in which several learning paradigms are combined to generalize well (Gupta et al., 2021). Future work may elaborate more on ensemble models or modular NLP pipelines particular to a sub-field of CTI (e.g. phishing, ransomware or nation-state APTs), whereupon optimization can be done at a finer grained level.

In terms of data, our analysis further shows how source diversity is an important element in developing CTI systems. A few threat intelligence providers (e.g., Talos, Kaspersky) have, in general, provided substantially more actionable material than others. This non-uniformity is consistent with the results of Farahmandian et al. (2021) that demonstrated that the quality and density of CTI highly vary among vendors. It also comes to light concerning the requirement of quality conscious data ingestion systems that may downgrade and prioritize the intelligence in respect of past content richness or source reputation.

Additionally, there is another crucial issue related to our entity linking evaluation, one more thing we have to take into consideration, and this is the interoperability with external knowledge bases. Although resources such as MISP and the NVD were found to give high accuracy in linking, others were not that consistent. This reflects the results of a study by Li et al. (2021) in which the authors have reported that knowledge bases may not be in sync with each other and have incongruence in the entity semantics. Therefore to Takeaway semantic normalization and entity reconciliation could be implemented in future CTI systems, either with graph neural networks or ontology alignment frameworks (Wang et al., 2019).

Deployment of CTI automation tools into actual operations is also a topic in which the considerations of scalability and real-time performance can significantly arise. Even though our BERT-large model provided the best accuracy, it is computationally-intensive, which can limit its use in latency-constrained domains, where Security Operations Centers (SOCs) are the best example. As Peng et al. (2021) observe, edge deployment might be considered with a lightweight model such as Distil BERT or Tiny BERT at an accuracy cost that is not critical. Additionally, new knowledge distillation and quantization techniques can be used to render these models more efficient (Jiao et al., 2020).

Lastly, ethical implications should be considered. Automation of CTI extraction evokes issues of bias, overfitting to vendor-specific language, and/or the spread of falsehood. Due to the harmful patterns that NLP systems trained on unverified or biased data can propagate, the latter can be especially detrimental when they are involved in automated defense processes (Vidgen et al., 2020). Such systems must, thus, have some human-in-the-loop validation mechanism, means of ongoing feedback integration, and a way of explaining model behavior.

Finally, this paper will present strong arguments as to why NLP, specifically transformer-based models, is effective in automating CTI extraction. Meanwhile, it emphasizes the lasting potential of hybrid designs, pre-constructed datasets, and post-validation of an architecture to establish real-world preparedness. Domain-adaptive training, cross-lingual, and real-time system engineering, and ethical governance of automation tools in the next generation of CTI tools should focus the research efforts in the future.

REFERENCES

1. Ahmed, A., et al. (2020). “A Survey on Cyber Threat Intelligence Integration with Machine Learning.” IEEE Access, 8, 217719–217740.

2. Akhtar, S., et al. (2021). “Cybersecurity Threat Detection Using NLP: Tools and Techniques.” IEEE Access, 9, 102146–102165.
3. Almukaynizi, M., et al. (2020). CTI sharing and STIX adoption: A case study. Computers & Security, 94, 101818.
4. Amoroso, E., Sledge, C., & Prasad, P. (2021). Deep learning for cyber threat intelligence: A comprehensive review. Journal of Cybersecurity and Privacy, 1(1), 41–66.
5. Ayoade, G., Khan, L., & Thuraisingham, B. (2018). Extracting cybersecurity entities from text using machine learning techniques. Journal of Information Security and Applications, 43, 76–85.
6. Bell, D., Cleland, G., & Millard, S. (2019). Ontology-based named entity recognition for cybersecurity texts. Knowledge-Based Systems, 165, 123–133.
7. Bird, S., Klein, E., & Loper, E. (2009). Natural Language Processing with Python. O'Reilly Media.
8. Bridge, R. A., et al. (2013). “Automated Threat Report Processing: Extracting Threat Indicators from Unstructured Text.” IEEE Security & Privacy, 11(6), 57–63.
9. Brown, T. B., et al. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901.
10. Chalkidis, I., Androutsopoulos, I., & Michos, A. (2020). Legal-BERT: The muppets straight out of law school. arXiv preprint, arXiv:2010.02559.
11. Chalkidis, I., et al. (2020). “Legal-BERT: The Muppets Straight Out of Law School.” arXiv preprint, arXiv:2010.02559.
12. Chen, Y., Zhou, S., & Xu, D. (2020). Domain-specific BERT models for cybersecurity applications. arXiv preprint, arXiv:2012.09852.
13. Cheng, Y., et al. (2020). “Detecting Obfuscated Malware Indicators Using Deep Contextual Embeddings.” Journal of Cybersecurity, 6(1), tna005.
14. Clark, K., et al. (2019). “What Does BERT Look at? An Analysis of BERT's Attention.” EMNLP.
15. Coppolino, L., et al. (2015). “Data Acquisition in Cybersecurity: A Real-Time Monitoring Framework for SIEM.” Computers & Security, 50, 199–211.
16. Deng, Y., Zhang, C., & Li, M. (2021). Few-shot learning for threat intelligence extraction using prototypical networks. IEEE Access, 9, 17429–17440.
17. Devlin, J., et al. (2018). “BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding.” arXiv preprint, arXiv:1810.04805.
18. Farahmandian, M., Dehghantanha, A., & Choo, K. K. R. (2021). Evaluating the quality of cyber threat intelligence feeds. Computers & Security, 102, 102127.
19. Farahmandian, M., et al. (2021). “Evaluating the Quality of Cyber Threat Intelligence Feeds.” Computers & Security, 102, 102127.
20. Ghafir, I., Hammoudeh, M., & Prenosil, V. (2021). Security monitoring and threat detection using big data: A review. Future Generation Computer Systems, 98, 591–605.
21. Gupta, A., et al. (2021). “Modular Transformers for NLP Tasks.” NeurIPS.

22. Husák, M., et al. (2018). “Survey of Attack Attribution in Computer Networks.” Computers & Security, 80, 1–15.

23. Jiang, L., Wang, H., & Xue, Y. (2018). Feature selection in text classification using conditional mutual information. Expert Systems with Applications, 89, 275–286.

24. Jiao, X., et al. (2020). “TinyBERT: Distilling BERT for Natural Language Understanding.” EMNLP.

25. Johnston, R., & Weiss, J. (2017). Cybersecurity for Industrial Control Systems. CRC Press.

26. Kumar, A., & Singh, S. (2020). “Automated Cybersecurity Text Mining Using NLP.” Journal of Network and Computer Applications, 167, 102738.

27. Kumar, A., et al. (2016). “Ask Me Anything: Dynamic Memory Networks for NLP.” ICML.

28. Kwon, D., et al. (2017). Identification of cyber threat actors using high-dimensional behavior vectors and machine learning. Digital Investigation, 23, 19–29.

29. Lee, J., et al. (2020). “Bio BERT: A Pre-trained Biomedical Language Representation Model for Biomedical Text Mining.” Bioinformatics, 36(4), 1234–1240.

30. Li, J., et al. (2021). “Entity Resolution in Cyber Threat Intelligence.” International Journal of Information Security, 20, 481–499.

31. Liao, Q., et al. (2016). “Automated Cyber Threat Intelligence Extraction from Text.” Proceedings of the ACM on Computer and Communications Security (CCS).

32. Lin, W., Duan, H., & Zhou, J. (2022). Fine-tuning domain-specific BERT models for CTI extraction. Cybersecurity, 5(1), 6.

33. Liu, Y., et al. (2020). “Fine-tune BERT for Extractive Summarization.” arXiv preprint, arXiv:1903.10318.

34. Ma, J., & Hovy, E. (2019). A self-attentive model with gated convolutional for named entity recognition. arXiv preprint, arXiv:1906.01477.

35. Marchetti, M., et al. (2017). “Filtering Network Traffic Data for Security Information and Event Management Systems.” Computers & Security, 66, 1–15.

36. Mittal, S., et al. (2019). “Cyber-All-Intel: An AI for Extracting Cybersecurity Threat Intelligence from Open Source Text.” Proceedings of the AAAI Conference on Artificial Intelligence.

37. Mohammad, S., & Somayaji, A. (2020). Automating cyber threat intelligence: A deep learning-based approach to extracting threat entities. Journal of Cybersecurity, 6(1), tna005.

38. Montemurro, M., et al. (2020). “Cybersecurity Knowledge Graphs for Threat Intelligence: A Review.” Computers & Security, 95, 101859.

39. Montemurro, M., et al. (2020). Cybersecurity knowledge graphs for threat intelligence: A review. Computers & Security, 95, 101859.

40. Pang, J., Zhang, C., & Liang, B. (2019). Information extraction for cybersecurity: A survey. ACM Computing Surveys, 52(6), 1–35.

41. Peng, H., et al. (2021). “Fine-tuned BERT for Cyber Threat Entity Extraction.” IEEE Transactions on Information Forensics and Security, 16, 1696–1705.

42. Peng, H., et al. (2021). “Knowledge Distillation and Compression of BERT Models.” ACM Transactions on Intelligent Systems and Technology, 12(3), 1–25.

43. Prior, A., et al. (2021). “Why Hybrid Models Work: Interpretability and Performance in Rule-Augmented NLP.” ACL Workshop on Responsible NLP.

44. Rasthofer, S., et al. (2017). “A Machine Learning Approach for Classifying and Categorizing Android Sources and Sinks.” NDSS Symposium.

45. Rastogi, M., Ghosh, A., & Yadav, S. (2021). Relation extraction techniques for cybersecurity threat intelligence: A survey. IEEE Transactions on Dependable and Secure Computing, 18(4), 1944–1961.

46. Rossi, R., et al. (2021). “Cybersecurity Threat Intelligence Text Mining Using Deep Learning.” Applied Sciences, 11(3), 1248.

47. Roth, D. (2017). “Incidental Supervision: Moving Beyond Supervised Learning.” AAAI.

48. Samtani, S., Chinn, R., & Chen, H. (2020). Cyber threat intelligence modeling using structured and unstructured sources. Information Systems Frontiers, 22(5), 1171–1187.

49. Tenney, I., et al. (2019). “BERT Rediscovers the Classical NLP Pipeline.” ACL.

50. Vidgen, B., et al. (2020). “Directions for Responsible NLP: Closing the Loop Between Ethical Development and Deployment.” ACL.

51. Wang, J., He, T., & Li, Z. (2021). Constructing cybersecurity knowledge graphs for enhanced threat intelligence. Computers & Security, 106, 102268.

52. Wang, Z., et al. (2019). “Knowledge Graph Alignment with Embedding Models.” IJCAI.

53. Yan, Z., et al. (2021). “A Span-Based Approach for Joint Entity and Relation Extraction.” ACL Findings.