



Review article

NLP-based techniques for Cyber Threat Intelligence

Marco Arazzi ^a, Dincy R. Arikkat ^b, Serena Nicolazzo ^c, Antonino Nocera ^a,
Rafidha Rehiman K.A. ^b, Vinod P. ^{b,d}, Mauro Conti ^d

^a Department of Electrical, Computer and Biomedical Engineering, University of Pavia, Italy

^b Department of Computer Applications, Cochin University of Science and Technology, India

^c Department of Science, Technology and Innovation, University of Eastern Piedmont, Italy

^d Department of Mathematics, University of Padua, Italy

ARTICLE INFO

Keywords:

Cyber threat intelligence
Natural language processing
Security
Named entity recognition
Knowledge graph
Large language model

ABSTRACT

In the digital era, threat actors employ sophisticated techniques for which, often, digital traces in the form of textual data are available. Cyber Threat Intelligence (CTI) is related to all the solutions inherent to data collection, processing, and analysis useful to understand a threat actor's targets and attack behavior. Currently, CTI is assuming an always more crucial role in identifying and mitigating threats and enabling proactive defense strategies. In this context, NLP, an artificial intelligence branch, has emerged as a powerful tool for enhancing threat intelligence capabilities. This survey paper provides a comprehensive overview of NLP-based techniques applied in the context of threat intelligence. It begins by describing the foundational definitions and principles of CTI as a major tool for safeguarding digital assets. It then undertakes a thorough examination of NLP-based techniques for CTI data crawling from Web sources, CTI data analysis, Relation Extraction from cybersecurity data, CTI sharing and collaboration, security threats of CTI, and role of LLM in this domain. Finally, the challenges and limitations of NLP in threat intelligence are exhaustively examined, including data quality issues and ethical considerations. This survey draws a complete framework and serves as a valuable resource for security professionals and researchers seeking to understand the state-of-the-art NLP-based threat intelligence techniques and their potential impact on cybersecurity.

Contents

1. Introduction	2
2. Related work	3
3. Methodology	3
3.1. Research objectives	4
3.2. Search strategy	4
3.3. Selection criteria	5
3.3.1. Inclusion criteria	5
3.3.2. Exclusion criteria	5
4. Overview of cyber threat intelligence	5
4.1. Threat intelligence life cycle	5
4.2. Types of cyber threat intelligence	5
5. Natural language processing in cyber threat intelligence	6
5.1. Preprocessing and tokenization	6
5.2. Text representation techniques	7
5.3. NLP algorithms and models	7
5.4. NLP libraries and tools	8
6. NLP-based techniques for threat intelligence collection	8
6.1. Crawling from clear web	9

* Corresponding author.

E-mail addresses: marco.arazzi01@universitadipavia.it (M. Arazzi), dincyrarikkat@cusat.ac.in (D. R. Arikkat), serena.nicolazzo@unimi.it (S. Nicolazzo), antonino.nocera@unipv.it (A. Nocera), rafidharehimanka@cusat.ac.in (Rafidha Rehiman K.A.), vinod.p@cusat.ac.in, vinod.puthuvath@unipd.it (V. P.), mauro.conti@unipd.it (M. Conti).

<https://doi.org/10.1016/j.cosrev.2025.100765>

Received 10 January 2024; Received in revised form 2 April 2025; Accepted 13 May 2025

1574-0137/© 2025 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

6.2.	Crawling from social media.....	10
6.3.	Crawling from dark/deep web.....	12
7.	NLP-based techniques for threat intelligence analysis.....	15
7.1.	Text classification.....	15
7.2.	Text similarity and clustering.....	18
7.3.	Text mining and open information extraction.....	18
7.4.	Cross-lingual threat intelligence.....	18
7.5.	Topic detection and trend analysis.....	19
7.6.	Text summarization.....	20
8.	NLP-based techniques for relation extraction.....	20
8.1.	Named entity recognition.....	20
8.1.1.	Annotation strategy.....	21
8.1.2.	NER approaches.....	21
8.1.3.	NER corpus.....	22
8.2.	Event identification.....	24
8.3.	Relation extraction.....	25
8.4.	Knowledge representation using graphs.....	29
8.5.	Explainability in CTI.....	32
9.	Cyber threat intelligence sharing and collaboration.....	33
9.1.	CTI information-sharing platforms.....	33
9.2.	Standardization efforts.....	34
9.3.	Quality and privacy in CTI sharing.....	35
10.	Adversarial attack.....	37
10.1.	Adversarial attacks in NLP.....	37
10.2.	Adversarial examples generation.....	37
10.3.	Attacks and defense techniques in CTI.....	37
11.	Role of large language models in CTI.....	38
12.	Challenges and open research directions.....	39
12.1.	Handling unstructured data.....	39
12.2.	Adversarial data injection.....	40
12.3.	Standardized quality assessment.....	40
12.4.	Privacy and ethical considerations.....	40
12.5.	Multilingual NLP for threat intelligence.....	40
13.	Summary and conclusion.....	40
	Declaration of competing interest.....	41
	Acknowledgments.....	41
	Appendix.....	41
	Data availability.....	42
	References.....	42

1. Introduction

Nowadays, Cyber Threat Intelligence (CTI, hereafter) has gained a paramount role in cybersecurity, providing organizations with valuable insights into the Tactics, Techniques, and Procedures (TTPs) employed by cyber adversaries. Indeed, since there has been a significant increase in the variety and number of cyber attacks and malware samples, security practitioners have started to rely on CTI to promptly recognize the indicators of a cyber attack, collect information about the attack methods, and respond to it accurately and timely. Hence, in broad terms, the CTI pipeline takes as input cybersecurity data and produces valuable insights aimed at enhancing proactive cybersecurity defenses. This includes formulating strategies to minimize the impact of cyber attacks and prevent them in the future.

Due to the intricate and multifaceted nature of the cyber security context, different organizations and experts may emphasize different aspects of CTI based on their specific needs, objectives, and contexts. For this reason, several definitions of CTI exist. For instance, one of the most popular definitions states that CTI is “evidence-based knowledge, including context, mechanisms, indicators, implications, and actionable advice about an existing or emerging menace or hazard to assets that can be used to inform decisions regarding the subject’s response to that menace or hazard” [1]. The authors of [2] state that CTI refers to “the set of data collected, assessed, and applied regarding security threats, threat actors, exploits, malware, vulnerabilities, and compromise indicators”. Finally, Dalziel [3] describes CTI as “data that has

been refined, analyzed, or processed such that it is relevant, actionable, and valuable”.

Despite the numerous CTI definitions and the various strategies for effectively utilizing them, a common baseline revolves around identifying suitable solutions to collect, manage, and analyze CTI data. In this context, recent research findings have proved that Artificial Intelligence (AI), and especially Natural Language Processing (NLP) techniques, play a crucial for several reasons. Firstly, because much of the information is represented by unstructured text data, such as threat reports, social media posts, news articles, and hacker forums, NLP allows CTI analysts to automatically handle these textual data and to enable the uncovering of relevant threats and trends. The analysis process of such information can be significantly sped up by automatically extracting key details from large volumes of text, such as Indicators of Compromises (IoCs) and TTP. Moreover, NLP can unravel the relationships between entities and events mentioned in text data, helping contextualize the information and building a more complete view of cyber threats and the actors behind them. Additionally, we can gain information about threat actors by analyzing online communications, forums, and other text sources through NLP techniques. Furthermore, because the CTI landscape is variegated, NLP can help to integrate information from various sources and formats into a coherent view, also assisting cybersecurity practitioners in the reports generation and sharing. Finally, NLP has also been exploited as a powerful tool to identify emerging threats and trends by analyzing large datasets of text-based information, thus proactively averting potential cyber attacks.

Taking all the above reasons into account, we can affirm that NLP plays a crucial role in enhancing the effectiveness and efficiency of CTI

by helping analysts cope with vast amounts of textual data, extracting actionable intelligence, and predicting the evolution of cyber threats. Nevertheless, a comprehensive survey on NLP-based techniques for CTI is still missing in the current scientific literature. Consequently, with this paper, we aim to fill this gap by providing a clear and updated picture of the CTI landscape, its multiple sources, and its gathering and sharing methods under a main perspective, focusing on NLP techniques and how they can enhance CTI ability to detect, analyze, and respond to cyber threats, effectively. In addition, in this paper we analyze the role of Large Language Models (LLMs) as a powerful tool for extracting knowledge from cyber threat sources.

We list our contributions as follows:

- We undertake a comprehensive examination of crawling techniques to extract data from main CTI Web sources, namely social networks, clear, and Dark Web.
- We analyze NLP models, techniques, and applications within the CTI domain. In particular, we give a detailed picture of the main papers dealing with text classification, similarity, clustering, summarization, cross-lingual, and topic detection in the context of CTI.
- We describe the most recent results of NLP-based techniques for relation extraction and representation.
- We address the topics of CTI standardization protocols and results sharing.
- We exhaustively examine the challenges and limitations of NLP in threat intelligence, including data quality issues, CTI security, and ethical considerations.

The outline of this paper is as follows. In Section 2, we present the related survey studies. Section 3 discusses the methodology we adapted to conduct this survey. Section 4 explains the CTI life cycle and different types of CTI. In Section 5, we make an overview of NLP techniques and their role in CTI. Section 6 deals with the different Web crawling techniques to collect CTI data from clear, Dark/Deep Web, and social networks. Section 7 details classical NLP-based techniques applied to CTI data, whereas NLP-based techniques for Relation Extraction are described in Section 8. Section 9 deals with papers about CTI sharing and collaboration. Section 10 focuses on security aspects related to CTI and, in particular, on adversarial attacks. In Section 12, we discuss several open challenges and insights for further research work. Finally, in Section 13, we draw our conclusion to the survey.

2. Related work

This section provides an overview of the existing literature surveys in the field of CTI.

Cascavilla et al. [4] focused on papers that deal with CTI data crawling, incident prediction and avoidance, and standardization of cyber-criminal activities in the Surface, Deep, and Dark Web. Similarly to our work, Rahman et al. [5] analyzed different extraction purposes, namely CTI text classification, IoCs and TTPs extractions. Moreover, they identified several types of textual sources for CTI extraction (i.e., hacker forums, threat reports, social media posts, and online news articles), and they observed that NLP and Machine Learning (ML) based techniques, such as supervised classification, Named Entity Recognition (NER), topic modeling, and dependency parsing, are the primary techniques used for CTI extraction.

The surveys presented in [6–8] analyzed the application of AI and ML in producing actionable CTI. While the authors of [6] only mentioned a work related to NLP and gave a brief survey on data collection and sharing. The authors in [7] conducted a short discussion on ML and AI tools (such as adversarial learning). The survey presented in [8] summarizes commonly used cyber security data sources with several approaches dealing with adversarial learning. The paper illustrated in [9] provides a systematic review of existing CTI platforms within the

industry; they reviewed the major data sources (without detailing the crawling techniques) and gave some insights on visualization, report generation, and intelligence dissemination. Moreover, they mentioned the enhancement of NLP and text mining as potential opportunities in the field of CTI.

The works described in [10,11] provide an analysis of the current role of the Dark Web as an environment that facilitates cybercrime and illicit gain. In particular, [10] compares state-of-the-art research studies that leverage the dark Web as an information source for CTI, describing their goals, approaches, tools, case studies, results, and possible limitations. A group of surveys focuses mainly on NER aspects related to the cybersecurity domain [12,13]. Specifically, the authors of [12] introduced common NER models, methods, and related resources.

The work presented in [14] focuses on Technical Threat Intelligence (TTI, hereafter) and the major problems related to it, CTI sharing, and evaluates the most known open source tools offering TTI. Surveys [15–18] specifically deal with CTI Sharing. In particular, El-Kosairy et al. [15] highlight the most recent contributions, which discussed how blockchain is integrated with CTI to solve the issue of CTI sharing. The survey presented in [16] gives an insight into diverse problems about CTI sharing between 2001 and 2018. Specifically, it deals with the way of establishing a threat-sharing program with decentralized stakeholders, focusing on what information can be shared, with whom, and how to automate some of the collaboration processes. Moreover, it also covers topics like CTI anonymization, encrypting the data, and presenting privacy risk scores. The work presented in [19] sheds light on existing standardization approaches for incident responses. Finally, Sun et al. [20] described state-of-the-art cybersecurity incident prediction schemes and methods; moreover, the type of datasets used in each work is identified and referenced in minute detail.

A survey on CTI mining techniques and the CTI knowledge acquisition taxonomy is presented in [21]. This analyzes the methodology that transforms cybersecurity-related information into evidence-based knowledge for proactive cybersecurity defense using CTI mining. Unlike our survey, it describes 53 papers published before 2022 and does not consider adversarial attacks. Moreover, we adopt a different classification scheme and deep dive into an extensive analysis of the data collection presenting different crawling techniques for gathering data from Clear, Dark/Deep Web, and Social Media.

Several recent contributions are presented in [22–24]. In particular, in [22] the authors analyzed current approaches and future directions regarding CTI sharing from 2011 to 2023. Arikkat et al. [23] evaluated Relation Extraction techniques in the CTI domain providing a picture of the main papers dealing with dependency-parsing, supervised and unsupervised learning approaches, and Large Language Models (LLM) from 2019 to 2024. Whereas the survey paper [24] focuses on LLMs utilized specifically to detect cyber threats limited to 74 papers.

Concerning the previous surveys reported in this section, our paper introduces important contributions to the research community. Indeed, first, we analyze the publications produced in a very recent period (i.e., from 2018 to 2025); this is a crucial aspect because, with the technological improvement of AI solutions, novel NLP-based strategies for CTI are more and more advanced, and effective to the results obtainable only a few years ago. Moreover, while existing works already consider some aspects that overlap with our survey, our proposal analyzes the recent CTI literature completely from the perspective of exploiting of NLP techniques. Thus it provides a comprehensive report of how NLP-based solutions have been successfully adopted in the reference context. Table 1 provides a summary analysis of the contributions of the existing related surveys compared to ours.

3. Methodology

In this study, we undertake a comprehensive literature review process, which can be outlined as follows: it involves establishing clear search objectives, entails the identification of relevant literature from

Table 1
Surveys related to our work.

Paper	Year	Literature timeline	Data Crawling	NLP for CTI	NLP for Relation Extraction	Adversarial Attack	Standardization Protocols	CTI Sharing
Abu et al. [17]	2018	2011–2017					✓	✓
Sun et al. [20]	2018	2013–2018	✓					
Tounsi et al. [14]	2018	2010–2017					✓	✓
Miloshevska [11]	2019	2015–2018						
Wagner et al. [16]	2019	2001–2018					✓	✓
Ibrahim et al. [7]	2020	2014–2020				✓		
Samtani et al. [9]	2020	2013–2019				✓		✓
Samtani et al. [8]	2020	2012–2020	✓					✓
Basheer et al. [10]	2021	2017–2021		✓				
Cascavilla et al. [4]	2021	2002–2020	✓	✓				✓
Gao et al. [12]	2021	2007–2019			✓			
Georgescu et al. [13]	2021	2007–2019			✓			
Montasari et al. [6]	2021	2015–2020	✓					✓
Sauerwein et al. [18]	2021	2012–2021						✓
Schlette [19]	2021	2001–2020					✓	
El-Kosairy et al. [15]	2023	2019–2022					✓	✓
Rahman et al. [5]	2023	2013–2022	✓	✓	✓			
Sun et al. [21]	2023	2014–2022		✓	✓		✓	✓
Alaeifar et al. [22]	2024	2011–2023						✓
Arikkat et al. [23]	2024	2019–2024			✓			
Chen et al. [24]	2024	2020–2023		✓	✓			
Our Survey	2025	2018–2025	✓	✓	✓	✓	✓	✓

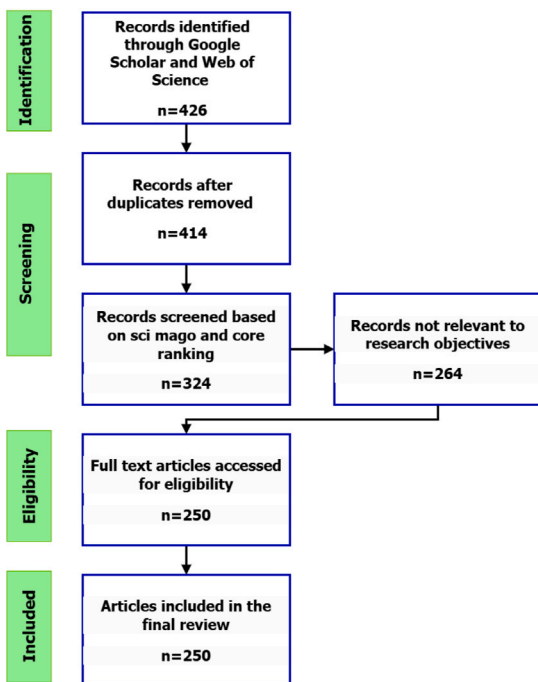


Fig. 1. PRISMA Flowchart for paper selection process.

journals and conference proceedings via a search engine driven by a well-defined search strategy, and encompasses the application of selection criteria to filter the articles during this iterative process. Fig. 1 presents the PRISMA flow diagram, providing a visual representation of the screening process. The diagram highlights the count of research works identified, excluded, and included.

3.1. Research objectives

The primary objective of this comprehensive literature survey is to conduct an in-depth analysis of NLP models, techniques, and their

applications within the CTI domain. This entails a detailed exploration of common data sources and the methodologies employed for data collection in the CTI field. The study delves deeply into the relevant literature, with a specific focus on papers that address topics like text classification, similarity, clustering, summarization, cross-lingual analysis, and topic detection, all within the context of CTI. The research extends its purview to encompass recent developments in NLP-based techniques for relation extraction. The survey also examines the literature related to adversarial attacks in the CTI field. Moreover, this research explores CTI standardization protocols and the sharing platforms, contributing to the enhancement of collaboration and knowledge dissemination in the CTI community. Furthermore, the study aims to identify areas in CTI where improvements are needed and to pinpoint potential research directions that can advance the capabilities of CTI. Lastly, this paper undertakes an exhaustive examination of the challenges and limitations that NLP encounters within the context of threat intelligence.

3.2. Search strategy

To gather relevant research focusing on NLP-based techniques for CTI, we designed a search strategy to align with our research objectives. We meticulously conducted searches through Google Scholar and Web of Science, thus accessing a wealth of scholarly resources. In addition, our search strategy encompassed a wide array of esteemed academic databases, such as the Institute of Electrical and Electronics Engineers (IEEE) Xplore, the Association for Computing Machinery Digital Library (ACM), ScienceDirect, and SpringerLink. The search scope encompasses a broad spectrum of publication years, spanning from 2018 to 2025, ensuring a well-rounded coverage of recent research in the field. Further, we developed the search terms to facilitate our initial exploration. The search terms assist in identifying relevant research and employ a combination of specific keywords and phrases to encompass various aspects of the threat intelligence and cybersecurity domain. The primary search term includes “threat intelligence”, “cyber intelligence”, and “CTI”. In conjunction with the primary term; we integrated additional search terms such as “indicators of compromise”, “tactics, techniques and procedures”, “hacker forum”, “dark web”, “APT attack”, “cyber attack”, “cyber security”, “named entity recognition”, “knowledge graph”, “cyber security relation extraction”, “adversarial attack”, “CTI platform”, and “CTI standardization”.

3.3. Selection criteria

In this section, we list the set of selection criteria we exploited to determine whether a scientific paper, identified through the search queries, is relevant to our study and reaches enough quality to be included in this survey. A paper is eligible for inclusion in the present work if it meets at least one of the inclusion criteria and none of the exclusion criteria applies. At the end of this screening, 192 papers have been selected based on inclusion/exclusion criteria.

3.3.1. Inclusion criteria

To evaluate the relevance of a paper and include it in our survey, we followed the subsequent criteria:

- we consider the corresponding author or the supervisor's importance in the field under analysis; including also non-peer-reviewed articles if they come from very important research groups or authors.
- we count the number of citations (we refer to Google Scholar [25] and Scopus; [26] Web sites to establish this value);
- we privilege more recent works;
- we take into account the importance of the journal (or the conference) where the paper has been published (we consider Scimago [27] and Core.edu [28] as ranking Web sites for journals and conferences respectively to establish this value).

3.3.2. Exclusion criteria

After the inclusion, the exclusion process is followed. A paper is excluded if only one of the following criteria is met:

- the paper is not written in English;
- the paper is not peer-reviewed and not coming from a very important research group or eminent author in the field or not presented at reputable conferences;
- the paper's year of publication is earlier than 2018;
- the paper is relevant in the NLP context but it is not specifically focused on CTI;
- the paper lacks relevance, it is an incremental refinement of an earlier proposed approach, there is a duplicate publication, or there is a most recent and cited advancement of the work published in a more important journal or conference.

4. Overview of cyber threat intelligence

CTI enables organizations to understand the evolving threat landscape, enhance their defenses, and respond effectively to emerging cyber attacks. Organizations adopt various types of CTI and go through various stages. The CTI life cycle and types are explained below. Table 2 illustrated all the acronyms used in the paper.

4.1. Threat intelligence life cycle

CTI is a data-driven process that helps organizations defend against cyber threats and attacks. It involves several stages, working together to collect, process, and analyze data, providing valuable insights for cyber security. This systematic process transforms data into actionable intelligence, requiring continuous adaptation to evolving needs and real-world events. Different researchers have highlighted various versions of the CTI life cycle, with Irshad et al. [29] mentioned six stages, Basheer et al. [10] and Ainslie et al. [30] mentioned four stages. They all typically follow a cyclical and iterative process, allowing for ongoing review and adaptation. The core principles of intelligence collection, analysis, and dissemination remain crucial in addressing the ever-changing cyber threat landscape. In light of existing literature, the main CTI phases can be grouped into six stages, namely Planning and Direction, Collection, Processing, Analysis, Dissemination, and

Feedback. The CTI process begins with *Planning and Direction*, where organizations establish their objectives, priorities, and requirements for threat intelligence. In the *Collection* phase, raw data is gathered from diverse sources, including open-source feeds, commercial threat intelligence providers, and internal logs. This data may include IoCs, malware samples, and network traffic logs collected through automated tools and manual research. Once data is collected, it needs to be refined and organized. *Processing* phase involves the conversion of raw data into a more structured format that can be used for analysis. In the *Analysis* phase, the processed data is analyzed to identify patterns, trends, and potential threats. Once the threat intelligence is analyzed and assessed, the findings are shared with relevant stakeholders during the *Dissemination* phase. Information is typically disseminated in a clear and actionable format, such as reports or alerts. Continuous improvement is a crucial aspect of a mature CTI system, with feedback collected from various stakeholders to evaluate the effectiveness of threat intelligence, identify areas for improvement, and address emerging requirements. The *Feedback* stage supports the refinement and adaptation of the CTI program to changing threats and organizational needs.

4.2. Types of cyber threat intelligence

CTI encompasses a wide range of subcategories, highlighting the diversity of information sources and purposes. These subcategories include open-source intelligence, technical intelligence, human intelligence, measurement and signature intelligence, social media intelligence, geospatial intelligence, signal intelligence, Deep or Dark Web intelligence, and communication intelligence [10]. However in most of studies [14,18], CTI classified into four types:

- *Strategic Intelligence*: Strategic threat intelligence aims to provide valuable insights for business decisions and processes, considering the impact of external factors such as geopolitics on adversaries' actions. Strategic threat intelligence relies on expert analysts and authoritative sources to collect and interpret data from various channels, facilitating organizations in gaining insights into past incidents, understanding the intentions of threat actors, recognizing emerging trends, and developing risk mitigation strategies. It encompasses aspects like financial consequences, attribution of attacks, and industry-specific perspectives. By connecting global events and policies with potential cyber threats, strategic threat intelligence empowers decision-makers to assess risks and align cybersecurity investments with their strategic objectives. This type of intelligence is typically conveyed through comprehensive reports. It provides top executives and IT management with the means to identify risks, threat actors, and the origins of breaches while focusing on long-term challenges and real-time alerts for critical assets.
- *Operational Intelligence*: Operational threat intelligence serves as a vital tool for organizations to gain a deeper understanding of potential threat actors, including their motives, capabilities, and opportunities for launching attacks, while also identifying vulnerable IT assets and assessing the potential consequences of successful attacks. This form of intelligence is typically amassed by government entities and greatly benefits incident response and forensic teams by enabling them to implement security measures for enhanced detection, early recognition of attacks, and safeguarding critical assets. The sources of operational threat intelligence predominantly comprise human intelligence, data from social media platforms, chat rooms, and real-world events linked to cyber attacks. It involves the analysis of human behavior and evaluating threat groups to forecast and prepare for forthcoming attacks. Typically presented in comprehensive reports, operational threat intelligence includes in-depth information about documented malicious activities, alerts about emerging threats, and suggested courses of action.

- **Tactical Intelligence:** Tactical threat intelligence is a pivotal component in the protection of an organization's assets, offering valuable insights into the TTPs employed by threat actors in cyber attacks. This intelligence category primarily serves the needs of cybersecurity professionals, including IT service managers, security operations managers, administrators, and architects. It empowers these experts with several capabilities such as gaining an in-depth understanding of adversaries' attack methodologies, anticipating potential data breaches, assessing the technical proficiencies and objectives of the attackers, and pinpointing the avenues of attack. Consequently, it enables security personnel to proactively devise strategies for detecting and mitigating threats, which may involve integrating known indicators into security products and rectifying system vulnerabilities. Tactical threat intelligence is sourced from various outlets, such as campaign reports, incident reports, malware analysis, attack group reports, and human intelligence. The collection of this data entails activities such as analyzing technical papers, collaborating with other organizations, and obtaining information from third-party sources. This intelligence encompasses intricate technical details, including the specifics of malware, details of campaigns, the methodologies employed in attacks, and information about the tools used, often documented in forensic reports.
- **Technical Intelligence:** Technical threat intelligence is centered on understanding the tools and resources used by attackers in their malicious activities, encompassing elements like command and control channels. In contrast to tactical threat intelligence, this form of intelligence has a shorter lifespan but is highly specific, enabling swift distribution and response to threats. While tactical threat intelligence covers the broader malware used in attacks, technical threat intelligence delves into the detailed implementation aspects of that malware, encompassing specifics such as IP addresses, domains, phishing email headers, and malware checksums. This intelligence is typically consumed by Security Operations Center (SOC) staff and incident response (IR) teams, which is crucial in identifying malicious activities. The indicators used in technical threat intelligence are sourced from active campaigns, attacks on other organizations, or data feeds from external parties and are often gathered during investigations of attacks on different entities. Security professionals leverage this information to bolster the detection capabilities of security systems like IDS/IPS, firewalls, and endpoint security systems. This, in turn, facilitates the identification of malicious traffic and suspected IP addresses involved in malware distribution and spam emails. Technical threat intelligence is directly integrated into digital security devices to block and identify malicious traffic, both incoming and outgoing, within an organization's network.

5. Natural language processing in cyber threat intelligence

This section is devoted to presenting the most common techniques employed in the NLP pipeline and how they can be applied to process CTI reports. In particular, we will describe the fundamental strategies of preprocessing and tokenization to sanitize the text from possible noise to enhance the robustness of the model built upon. Following that, we will introduce different representation techniques, from the most basic (in Section 5.2) to the most advanced ones, that make use of Deep Learning (DL) models to obtain more sophisticated and contextualized representations (in Section 5.3). Then, we will introduce the most modern NLP models that represent the state-of-the-art in combination with the tool that implements them.

5.1. Preprocessing and tokenization

Preprocessing and tokenization are fundamental preliminary cleaning phases in the NLP task pipeline. Data cleaning is important to

Table 2
Summary of the acronyms used in the paper.

Symbol	Description
APT	Advanced Persistent Threat
CTA	Cyber Threat Actor
CTI	Cyber Threat Intelligence
CVE	Common Vulnerabilities and Exposures
DNF	Dark Net Forum
DNM	Dark Net Marketplace
FinTech	Financial Technology
GNN	Graph Neural Network
IDS	Intrusion Detection System
IoC	Indicator of Compromise
IRC	Internet Relay Chat
KG	Knowledge Graph
ML	Machine Learning
NER	Named Entity Recognition
NIST	National Institute of Standards and Technology
NLP	Natural Language Processing
OIE	Open Information Extraction
OSINT	Open Source Intelligence
OSN	Online Social Network
SVM	Support Vector Machine
TTI	Tactical Threat Intelligence
TTPs	Tactics, Techniques & Procedures
CRF	Conditional Random Field
STIX	Structured Threat Information Expression
BERT	Bidirectional Encoder Representations from Transformers
XAI	eXplainable Artificial Intelligence
NVD	National Vulnerability Database

reduce the dimension of each text, remove noise, and improve the performance of models that take sanitized text as input. Among the common strategies to preprocess the data [31] there are (i) removing HTML tags if the crawled resource is an HTML page; (ii) replacing not alphanumeric characters; (iii) converting text to lowercase; (iv) removing stop-words, duplicates, common words (i.e., appearing in more than 85% of the documents), and rare words (i.e., appearing fewer than three times in the entire corpus). A common tool for data preprocessing is Stanford CoreNLP [32]. Also, the studies that leverage social network posts (e.g., tweets) follow similar preprocessing steps. For some of these steps (i.e., removing stop words and stem tokens) [33] utilizes the functionalities available in the Natural Language Toolkit (NLTK) libraries [34].

To disambiguate terms, which are inflected forms of the same word, two techniques are the most common: lemmatization and stemming. The former consists of substituting the inflected forms in the text with the original word and the latter, instead, reduces the inflected token to their root. These techniques have an important role also in the CTI approaches that involve NLP. An example of this can be seen in [33,35–41] where the authors apply NLP-based strategies to extract CTI from unstructured sources, automatically. In particular, they applied a sanitization phase that includes stemming, lemmatization, and stopword/tag removal, to the descriptions of adversarial techniques. Models like BERT [42] require a sequence of a fixed length as input. In these situations, the input must be preprocessed by either truncating it or adding a padding sequence to reach the required length and convert input tokens into numerical identifiers. For example, this approach is adopted by the authors of [43–46] where they use BERT to perform Named Entity Recognition (NER) and threat classification in CTI. In particular, in addition to padding and truncation, they took into consideration a case-sensitive tokenizer because capitalized letters could be informative of phrases referring to a ransomware name. The authors of these papers also use tokenizers to split complex words into terms to make them identifiable in a reference vocabulary. Table 3 reports a summarized view of the different techniques adopted by reference papers in the CTI domain.

Table 3
Preprocessing and Tokenization.

Paper	Year	Lemmatization/ Stemming	Stopwords Removal	Tokenization
Ji et al. [40]	2019	✓	✓	✓
Dasgupta et al. [45]	2020	×	×	✓
Kadoguchi et al. [38]	2020	✓	✓	✓
Yang et al. [39]	2020	✓	×	✓
Evangelatos et al. [43]	2021	×	×	✓
Gao et al. [36]	2021	✓	×	×
Gao et al. [37]	2021	✓	×	×
Orbinato et al. [35]	2022	✓	✓	✓
Kuehn et al. [44]	2023	×	×	✓
Wang et al. [46]	2023	×	×	✓

5.2. Text representation techniques

As already mentioned in the previous section, neural networks and statistical algorithms cannot interpret text without converting it into a numerical representation. Representation techniques can leverage different strategies, ranging from statistical approaches, such as the very common Bag-of-Words [47], or representation obtained by a neural network, as done by Word2Vec [48].

As for statistical approaches, Bag-of-Words is a well-known strategy that represents text by counting the occurrences of the words in the corpus of the documents. The idea behind this approach is to measure the importance of a word according to its frequency about a specific outcome of the target ML algorithm. With this strategy, word frequencies are considered with the same relevance in the computation of the final result. However, some of the most frequent words are noise. To solve this problem, some enhancements to the Bag-of-Words strategy have been proposed. One of the most common solutions in this sense is TF-IDF [49]. TF-IDF simply counts the word occurrences and weights their contributions according to how frequent they are in the different documents.

Another method based on statistical analysis is GloVe [50]. GloVe is an algorithm that generates vector representations of words, building the co-occurrence matrix between words and encoding the information regarding the probability ratio of co-occurrence probabilities of words. As for the approaches using neural network representations, a well-established method is Word2Vec. Unlike previous methods, Word2Vec uses a neural network to create a space where words are represented as vectors, and their positions in this space reflect how they relate to each other. The Word2Vec approach proposes two strategies. The first strategy is called Continuous Bag-of-Words (CBOW) and is used to generate word representation starting from the representation of the context. The second one, called Skip-Gram, is used to generate the context vector starting from a single word.

The strategies introduced above are fundamental tools to generate representations of text to be fed into NLP models targeted to CTI tasks like, topic modeling or classification [32,38,51–61]. For instance, the solution proposed by [51] applies a TF-IDF strategy to generate representations of malicious URLs that are then fed to a K-means algorithm to obtain the different topics.

In [55] Bag-of-Characters representation, a Bag-of-Word strategy applied at a character level is used to generate the embedding of short-text CTI reports to be processed by a BiLSTM network for the NER task. Similarly, in [54], Word2Vec and GloVe are used in the same way on short text documents about CTI from Twitter. Another example of this is proposed by [38], where the authors analyze posts from the Dark Web. To do so, they exploited *doc2Vec*, which works analogously to Word2Vec but instead of generating word representations it produces vector embeddings of the entire document. Then, the obtained vectors are used as input to an unsupervised clustering algorithm. The approach presented in [52] uses an embedding matrix generated using GloVe as input to a C-BiLSTM network that aims

Table 4
Text representation techniques in the literature.

Paper	Year	Bag-of-Words/ Bag-of-Characters	TF-IDF	GloVe	Word2vec/ doc2Vec
Ayoade et al. [62]	2018	×	✓	×	×
Goseva et al. [63]	2018	×	✓	×	×
Bo et al. [58]	2019	×	✓	×	×
Dionisio et al. [54]	2019	×	×	✓	✓
Kadoguchi et al. [53]	2019	×	×	×	✓
Le et al. [56]	2019	×	✓	×	×
Li et al. [64]	2019	×	✓	×	×
Nayak et al. [51]	2019	×	✓	×	×
Queiroz et al. [65]	2019	×	×	×	✓
Adewopo et al. [57]	2020	×	✓	×	×
Ampel et al. [52]	2020	×	×	✓	×
Kadoguchi et al. [38]	2020	×	×	×	✓
Kim et al. [55]	2020	✓	×	×	×
Kristiansen et al. [59]	2020	×	✓	×	×
Li et al. [61]	2020	×	✓	×	×
Zhao et al. [32]	2020	×	×	×	✓
Pantelis et al. [60]	2021	×	✓	×	×
Biswas et al. [66]	2022	×	✓	×	×

at exploiting source code labeling. Analogously, the work presented in [53] generates document representations using *doc2Vec*, as done by [38], but in this case, it is used to classify cyber attacks using a Multilayer Perceptron (MLP) network.

A schematic report of the techniques used by the papers analyzed in this section is visible in Table 4.

5.3. NLP algorithms and models

In Section 5.2, we presented well-known and established representation techniques. Nowadays, more advanced neural models have emerged to build sophisticated techniques of representation and learning. One of the main drawbacks of previous approaches is the inability to have different representations of the words in different contexts. For instance, algorithms such as Word2Vec and GloVe generate the same word representations in any context. To overcome this flaw, modern models like ELMo [67] or those based on the transformer architecture can generate different embeddings according to the surroundings of the words [42]. ELMo uses a bidirectional LSTM trained to generate different contextualized representations for each word to solve the ambiguity of words with different meanings. In particular, the ELMo model has been trained on a large set of data to predict the next word of a sequence. With such characteristics, this model is capable of generating generally enough representations to be used in different tasks or even improved with accurate fine-tuning.

This aspect is shared also with transformer-based models like BERT [42] and GPT [68]. The original architecture of BERT is characterized by an encoder-only model proposed in two versions that differ in the number and dimension of their encoding blocks. The basic one is composed of 12 encoding blocks with 768 hidden units each; the larger model is characterized by 16 blocks with a size of 1024 units. Compared to LSTM models, transformers are unaware of the concept of sequence. To solve this issue, a positional embedding layer is added at the entrance of the BERT model to encode in each word representation information about their position in the sentence. One of the main characteristics of BERT refers to the strategy used to pre-train the model. The training has been performed on two different tasks contemporaneously. In particular, in the first task, the model objective is to predict the original value of a masked word in a sentence. With the second task, instead, the authors train BERT to predict the next sentence. Like ELMo, BERT can be used to produce different representations of a word according to the context. The second example of a transformer-based model is GPT, which recently received huge attention due to its text generation capabilities [69–71]. Different from BERT, GPT is characterized by a decoder-only architecture. The authors

Table 5
NLP solutions used for Text Representation in the reference literature.

Paper	Year	ELMo	BERT	GPT
Yin et al. [76]	2020	✗	✓	✗
Ranade et al. [80]	2021	✗	✓	✗
Ranade et al. [78]	2021	✗	✗	✓
Setianto et al. [79]	2021	✗	✗	✓
Alam et al. [75]	2022	✗	✓	✗
Chan et al. [41]	2022	✗	✓	✗
Liu et al. [77]	2022	✗	✓	✗
Wang et al. [46]	2023	✗	✓	✗
Zhou et al. [72]	2023	✓	✓	✗
Bayer et al. [73]	2024	✗	✓	✗
Ferrag et al. [74]	2024	✗	✓	✗

proposed different evaluations of their GPT model over the years. In its latest version, GPT-4, the training is characterized by a reinforcement learning phase that uses rewards given by humans. This strategy allows the model to achieve extra quality in the produced text compared to the previous versions.

As said earlier, due to the large amount of data used for the training, the models discussed in this section are general enough to be used as a backbone in NLP pipelines in the CTI field. This is demonstrated by the authors of [72], who test their proposed NER dataset with different combinations of models that exploit representation produced by ELMo and BERT. Another example can be seen in [46], where the authors slightly modify the input of a BERT model to adapt it to a different objective, i.e., restoring the original order of a permuted text. Similarly, two recent works [73,74] describe pre-trained, general-purpose cybersecurity language models based on BERT which are CySecBERT and SecurityBERT, respectively. In particular, SecurityBERT is specifically intended for cyber threat detection in IoT networks. The authors of [75] proposed a Python library for NER tasks on CTI documents, which allows the import of pre-trained BERT models. In [76], the authors presented EXBERT, a framework that uses BERT as a backbone to generate the embeddings that are hence provided as input to a classifier. The objective is to predict the exploitability of vulnerabilities from their software descriptions. The approach proposed by [77] uses BERT for data augmentation in the context of cybersecurity. In [78], a GPT-2 fine-tuned model is used to generate fake CTI reports to demonstrate a data poisoning attack on a knowledge extraction system. Instead, the authors of [79] proposed to use a GPT-2 fine-tuned model to parse Unix commands from log files in real-time. A schematic summary of the NLP tools used by the papers described above is visible in Table 5.

5.4. NLP libraries and tools

The processing procedures presented in 5.1 are well-known and established techniques implemented in many libraries. Two of the most popular choices are the libraries implemented in Python NLTK and spaCy. NLTK is a very easy-to-use library that implements many different useful tasks for preprocessing and also basic algorithms for clustering, such as K-means and classification, e.g., Naive Bayes. As for the preprocessing tasks, it implements several lemmatizations and stemming approaches and provides lists of stopwords to remove in different languages. SpaCy, instead, provides more advanced tools like wrappers for neural network models implemented in the main frameworks Tensorflow and Pytorch with GPU support. Interestingly, in the latest version of spaCy at the time of writing this article, developers added support to transformer-based models. Another powerful tool that offers advanced capabilities like tokenization, POS-Tagging, and NER is the one proposed by Stanford, called Stanford CoreNLP, available for both Python and Java languages. When it comes to transformer-based models, such as BERT and GPT described in Section 5.3, one of the most well-known and complete ecosystems providing easy and intuitive access to such models in an NLP workflow is Huggingface. Huggingface

allows an easy import of pre-trained models and associate tokenizers to be used for many different tasks, including fine-tuning. In addition, the developers offer an open hub where it is possible to upload fine-tuned models freely downloadable by the community.

Most of the papers in the literature that involve NLP [33,35,56,60,63,81,82] use the NLTK library to exploit implemented algorithms, such as TF-IDF vectorizer, stop-words or stemming/lemmatization. However, many others [36,37,40,77,80,83] use spaCy for vector representation and stemming/lemmatization. Stanford CoreNLP is typically used as an additional advanced library by many solution-proposing approaches for NER and POS-tagging [43,46,54,83,84]. Finally, as we stated before, transformer-based models are recently becoming very popular, and pre-trained ones are often exploited thanks to the support of the Huggingface community [41,43,80]. In such papers, pre-trained transformer-based models are used after a fine-tuning step for the specific task of CTI reports analysis.

Lessons learned: From this section we can observe that, as a significant portion of threat intelligence data is unstructured, researchers employ NLP techniques such as stop-word removal, lemmatization, stemming, and word vectorization using techniques like TF-IDF, Word2Vec, Doc2Vec, GloVe, and BERT. Traditional methods such as stop-word removal, lemmatization, stemming, and TF-IDF help preprocess textual data by reducing noise and improving relevance. However, TF-IDF lacks contextual understanding. Word embedding techniques like Word2Vec, Doc2Vec, and GloVe improve upon this by capturing semantic relationships, helping to detect attack patterns and emerging threats. However, they assign fixed meanings to words, which can lead to misinterpretations in cybersecurity contexts where terms may have multiple meanings (e.g., “worm” referring to malware vs. a biological organism). Advanced transformer-based models such as BERT, SecurityBERT, and CySecBERT address these limitations by providing context-aware word representations, significantly improving entity recognition and relation extraction in CTI. SecurityBERT and CySecBERT, specifically trained on cybersecurity texts, outperform general NLP models in identifying IoCs, malware names, attack techniques, etc. Although, BERT-based models are better at understanding security-related texts, a hybrid approach that combines traditional NLP techniques for initial filtering with deep learning models for contextual analysis can enhance CTI extraction across diverse data sources

6. NLP-based techniques for threat intelligence collection

This section aims to survey the solutions for CTI data gathering from Online Social Networks (OSNs, for short), the Clear Web, and the Dark/Deep Web, where malicious actors collaborate and communicate to plan cyber attacks. Not all information gets published into standard CTI databases and appliances; CTI is often shared in unstructured ways like blogs, posts, or threat reports from security companies or experts. For this reason, multiple online data sources are used as signals to generate warnings indicative of new potential cyber threats. Information gathering is the first and crucial step for collecting relevant data about new vulnerabilities, exploits, security alerts, threat intelligence reports, and security tool configurations. From the collection perspective, data can be divided into two categories: (i) indicator-based data, which mainly provides IoCs for quick attack prevention, and (ii) document-based data, which may contain richer and more comprehensive threat information than the former category, but, at the same time, may require more complex NLP techniques for the analysis. The first group is, in turn, divided into high-level IoCs (represented by TTPs, malware, and tools) and low-level IoCs (e.g., IP, URL, hash, domain name, source/destination port, timestamp, and infection type). In both cases, preliminary steps for crawling Web pages about malicious content involve (i) gathering URLs on the internet, (ii) filtering out benign content, and then (iii) downloading malicious ones.

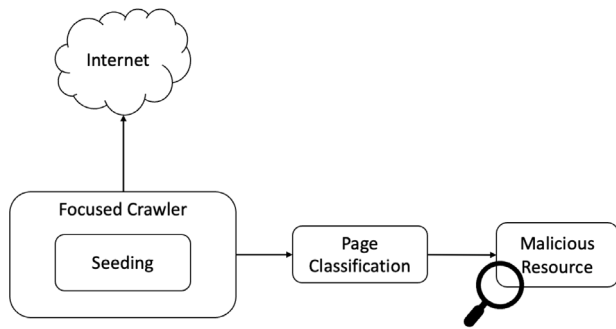


Fig. 2. General Architecture for crawling the Clear Web.

6.1. Crawling from clear web

The clear (also Open or Surface) Web represents the standard, publicly accessible portion of the Internet, where information is openly available, indexed by search engines, and accessible via conventional Web browsers. Since their introduction, Web crawlers have been intended to gather data from Clear Web. In Fig. 2, a general architecture for crawling this kind of Web is shown. Focused crawlers scan the Internet starting from some seed URLs. Usually, a *Seeding* module is in charge of locating initial URLs to be used by focused crawlers. Since they only search for relevant information, their performance can go down if different clusters of Web pages are isolated; hence, choosing a high number of appropriate seeds can solve this issue. To improve the performance of the overall architecture, inside the Seeding component, ML models can be used to classify URLs before navigating the Web pages. Finally, a page classification module is launched to assess if the Web page includes malicious content.

In [85], the authors proposed a two-stage architecture, including a crawling module and a content ranking module. The first part produces a focused crawler that employs NLP to decide if some target Web sites deal with specific cybersecurity-related vulnerabilities. Specifically, a Support Vector Machine (SVM) classifier is combined with a user-provided query relevant to the topic. Hence, concerning the reference architecture in Fig. 2, the Seeding module is enriched with this facility to improve seeds' selection. Moreover, the second component is in charge of assessing the relevance and usefulness of the crawled content. In the extended version of the paper [86], the authors presented an architecture called inTIME, including a more complex crawling infrastructure built on top of NYU's ACHE crawler and enhanced through the use of regex-based filters to direct the crawl to specific parts of a Web site. Moreover, to support focused crawling, inTIME includes the SMILE page classifier, which is based on an ML-based text classifier (e.g., SVM and random forest), trained by a selection of positive and negative examples of Web pages, to direct the crawl toward topically relevant Web sites.

The issue of discriminating between benign and malicious Web pages before crawling their content is also faced by "MalCrawler", which has been designed to crawl and search malicious Web sites efficiently [87]. In this system, the set of URLs representing the initial seeds' collection is processed, and only malicious links are navigated. MalCrawler can also prevent redirections and cloaking, generally used to avoid showing pages containing malicious content to search engine crawlers. The libraries used to develop MalCrawler are: (i) JSoup Library for parsing Web pages and extracting hyperlinks, document content, and JavaScript tags; (ii) Rhino a JavaScript Emulation Library to analyze the runtime behavior of JavaScript; and (iii) HTML Unit a browser emulation library written in Java to test redirection and cloaking. Another recently presented focused crawler, called ThreatCrawl [44], leverages the aforementioned general architecture, and it is based on both a retrieval and an extraction step. Additionally,

this system presents a monitor module as well, that stores the state of each retriever and extractor thread and can check for any critical error. Moreover, to extract the main content of an HTML Web Page and return it as parsed text, ThreatCrawl leverages both Trafilatura [88] and BeautifulSoup. Both the works presented in [84,89] design a cybersecurity engine, according to which the data collection step is done manually from multiple data sources, which are updated periodically via a monitor. Specifically, the main scope of [84] is to model and visualize security datasets based on users' queries and demands. Three datasets are exploited in this work, namely: (i) 1000 security news articles [90] mentioning security events and annotated by experts; (ii) some vulnerability archives collected from authoritative vulnerability database; (iii) several tweets related to security keywords and extracted using a Python library called Twitterscraper.

Similarly to the previous paper, also in the work presented in [91], the authors crawl multiple data sources. To do so, a semantic search system is developed to facilitate searching for high-level IoCs in the CTI corpus. They identified 36 Cyber Threat Actors (CTAs) and collected publicly available documents describing both the attack incidents (APT) and the CTAs. As also done in [92–94], they leveraged a publicly available repository called APTNotes containing cyber threat reports since 2008. Using the search engine and the repository, they collected 327 unstructured CTI documents. The work described in [94] collects a set of publicly available reports starting from MITRE ATT&CK (a globally accessible knowledge base of adversary tactics and techniques based on real-world observations); and APT Groups and Operations, which link threat actors to relevant reports. Through a manual expert review, the authors selected the most important threat actors for which there are a sufficient number of reports. The final dataset contains 249 reports, which describe attacks performed by 12 different actors. Similarly, the authors of [95] leveraged a dataset of corpora collected on the XianZhi platform. Recently, Arikkat et al. [96] compiles an extensive list of reliable Web sites renowned for publishing security-related blogs from which they extracted 39,166 blog articles via a multi-page Web Crawler leveraging Python libraries like BeautifulSoup and Selenium.

Another important source of CTI data can be found in the authoritative cybersecurity databases available online. Authors of several papers [40,82,84,97–99] leveraged data from online databases. Often, they used this data in combination with posts gathered from the Twitter Social Network as a ground truth database. The most popular employed databases are listed in the following:

- Common Vulnerabilities and Exposures (CVE) database. Each vulnerability present in this storage is identified by a CVE ID. The Web portal allows the user to make queries by CVE IDs to gather detailed information about the already known threats as well as to provide a Common Vulnerability Score (CVS) evaluating the impact of the given vulnerability.
- Common weakness enumeration (CWE) is a community-driven project that aims at identifying and classifying common software security weaknesses in a standardized and systematic way. Moreover, for each weakness listed in the CWE, a piece of associated information is present about how to identify, mitigate, and prevent the vulnerability.
- Common Attack Pattern Enumeration and Classification (CAPEC) database provides a broad hold of a wide list of attack patterns, which are widely used by security analysts, developers, and testers. Similarly to CWE, CAPEC categorizes attack patterns into a hierarchical structure. Each attack pattern is assigned to a unique identifier and organized into different categories and families.
- The Web Application Security Consortium (WASC) collects and organizes the threats related to the security of Web sites. Its main focus is to categorize and classify various Web application security threats and vulnerabilities, making it easier for organizations to understand and address them.

Table 6
Crawling from Clear Web.

Paper	Year	Technology	Seeding	Type of CTI extracted	Sample Size	Feasibility
Singh et al. [87]	2017	JSoup Library, Rhino, HTML Unit	From “Malware Domain List” ^a Website	Browser components (like, plug-ins); Attribute and Parameter Values in Method Calls; Sequence of Execution of Method Calls	A not defined number of Webpages	✓
Li et al. [89]	2018	Manual Collection	–	CVEs, victimized devices, device manufacturers and impacted locations	610 collected articles	✓
Dong et al. [82]	2019	CVE and NVD databases	–	Software name and version	78,296 CVE; 78,296 NVD entries	✓
Ji et al. [40]	2019	PRC and Hackmageddon databases	–	Event type, date, victim organization(s), and short description	61,957 cybersecurity events	✓
Koloveas et al. [85]	2019	ACHE open-source focused crawler	SVM classifier and keywords	User posts, comments, and tags	20,000 Web sites	✓
Noor et al. [91]	2019	Manual Collection	36 cyber threat actors identified by the authors	Attack observables, indicators, Tactics, Techniques and Procedures (TTP), incidents, threat actors, campaigns, exploit targets and course of actions	327 unstructured documents	✓
Perry et al. [94]	2019	Publicly available reports	–	Threat actors and relevant text reports	249 reports	✓
Alves et al. [98]	2020	CVE, PacketStorm, and other CTI databases	–	CVE-ID, a short description, and the creation date, analysis of the security issues raised by the vulnerabilities, known exploits, and possible fixes or mitigation actions	455,026 entries	✓
Wu et al. [95]	2021	XianZhi platform	34 hand-crafted TTPs	Text corpus, webpage URL and publish time of the Webpage	971 Web sites	✓
Sun et al. [84]	2022	Manual Collection	–	Text corpus, link, author, date, comments, title	1000 news articles [90] and vulnerability archives from authoritative vulnerability database	✓
Vishnu et al. [99]	2022	CVE database	–	CVE-ID, description, CVE type, date	134,091 vulnerability descriptions	✓
Kuehn et al. [44]	2023	Trafilatura and BeautifulSoup	List of URLs	Web page, URLs, TTPs	259 URLs	✓
Arikkat et al. [96]	2024	Multi-page Python Web Crawler	20 sources identified by the authors	Web page text, title and content.	39,166 blog articles	✓

^a <https://precisionsec.com/threat-intelligence-feeds/malware-domain-list/>.

- NIST’s National Vulnerability Database (NVD) mirrors and complements CVE entries on their database. It is maintained by the National Institute of Standards and Technology (NIST), which is a federal agency of the United States Department of Commerce. Every hour, NVD contacts CVE to obtain new CVE IDs of recently disclosed vulnerabilities.
- PacketStorm is an online resource and repository for a variety of current and historical cybersecurity-related information, tools, and databases.
- Privacy Rights Clearinghouse (PRC) is an independently maintained collection of reports about cybersecurity incidents divided according to victim organization type (i.e., government agencies, businesses, medical service providers, educational institutions, etc.).
- Hackmageddon is another reputable collection of public reports about cybersecurity incidents.

Table 6 summarizes, for each analyzed paper, the technology used to crawl data, the strategy seeding has been performed, the type of CTI extracted, the amount of data crawled, and the feasibility of the method.

6.2. Crawling from social media

Online social networks provide an open environment where both offensive and defensive practitioners can engage in discussions, report

incidents, and promote timely information about vulnerabilities, attacks, and malware. Among the online social networks, Twitter has always been one of the most popular in the scientific research context, mainly because of (i) its large and international user base (almost 400 million users), (ii) its set of available APIs, that allows developers to easily integrate Twitter data into their systems, and (iii) the tagging functionality that allows users to label their tweets and search them via keywords (i.e., hashtags), thus providing a natural data aggregation capability. For all these reasons, Twitter has become the preferred platform for disseminating up-to-date IoCs and threat intelligence data and conducting threat hunting [100,101].

One of the main issues of gathering data from Twitter is to wisely choose starting keywords to filter the stream listener results and reduce the amount of irrelevant information that is gathered. Nevertheless, crawling from this social network is quite a standard task for all the papers considered in this section and mainly proceeds with the execution of three tasks: data collection, keyword-based pre-filtering, and text pre-processing.

- In the *collection* phase, an initial set of Twitter accounts is required. Starting from this list, the crawler usually queries Twitter’s streaming API for tweets.
- In the *pre-filtering* phase, a set of user-defined keywords are used to drop irrelevant tweets.
- In the *pre-processing* phase, some transformation on data is performed to make representation uniform.

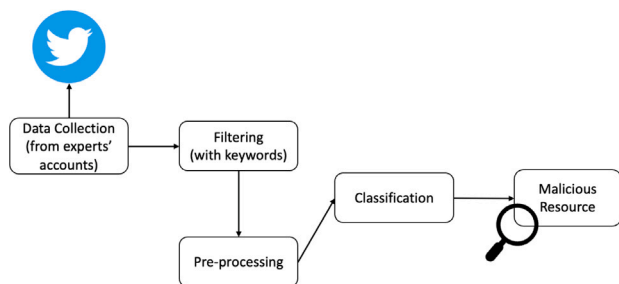


Fig. 3. General architecture for crawling data from Twitter.

Fig. 3 summarizes the main steps explained above.

The works presented in [54,57,102–105] follow this standard crawling architecture to crawl tweets and classify them as potential cyber threats. As for the collection step, in [103], the authors identified accounts related to Hacktivists, cybersecurity feeds, security researchers, and companies. Then they crawled tweets from their timeline, and after that, they filtered these tweets again based on a security keyword filter and a keyword list obtained by the Department for Homeland Security [106]. Also, the framework IoCMiner [107] starts from an initial set of Twitter accounts to collect tweets about CTI. This initial set is built through Twitter user-defined lists of identified CTI experts with a mechanism of validation that measures the relevancy, popularity, and comprehensiveness of the list in addition to the credibility of its owner. Moreover, it employs a CTI classifier to further filter out non-CTI tweets from the observed data streams. Finally, IoCMiner uses a set of regular expression rules to extract IoCs from the identified tweets.

Similarly, the framework presented in [104], called DISCOVER, aims to process data from multiple sources (such as Twitter and security blogs), and by employing data mining techniques, it identifies novel terms related to a potential cyber threat in the form of a warning alarm. To gather data from Twitter, they leveraged some experts' profiles, which it identifies in the timeline of international researchers and security analysts. As for the security blogs, the authors manually chose a list of 290 blogs and extracted data from them via a custom crawler. Data from these two sources are filtered to find words related to cyber threats.

The works analyzed up to this point collect raw data from experts' profiles and then apply a filter based on keywords on their timelines. The framework called #Twiti [108], instead, performs data collection in two ways: user tracking and keyword tracking. To perform the latter modality, it extracts the top 100 words that appear in tweets containing IoCs and uses them as starting keywords. Moreover, the presence of external sources is detected and also used as search criteria (such as URLs to security vendor blogs or Pastebin.com).

In [33,105,109], the authors presented a framework for the detection and classification of cyber threat indicators in the Twitter stream. Specifically, for [33] the data collected consists of a corpus of 21,000 tweets gathered with a custom stream listener developed through Tweepy. This work proceeds only by a list of keywords as a pre-filter for the stream listener, and after that, it performs two classification steps assessing the relevance of the cyber threats and the threat types. The first step is performed by leveraging the topic modeling API of IBM's Watson Natural Language Understanding service [110]. [109], instead, is done manually by simple string matching on the rough tweet text to find the type of thread. The work presented in [105] design an API's main component that is a crawling using the Twitter API's streaming feature to continuously extract tweets containing IOCs. Similarly to [33], the framework CTI-Twitter [59] uses a set of manually decided keywords as input for the streaming API, but it also leverages a Python library known as GetOldTweets to enrich this set with a historical collection of tweets. In [109], authors presented a work to identify target domain-relevant feeds containing critical patterns, and

they preprocessed the whole tweet stream, filtering it via keywords. The chosen keywords are related to traditional cybersecurity events, including account hijacking, data breaches, and denial of service attacks, and also to emerging events, such as ransomware and cryptocurrency mining malware. Interestingly, the authors also applied dynamic query expansion (DQE) to enrich tweet collection with tweets semantically related to cybersecurity.

The dataset used by [97] comprises tweets collected using Twitter API and information crawled from vulnerability information databases, such as NVD. They linked the data from the two sources together if they refer to the same CVE ID. In particular, the initial seeds' set is composed of searching on Twitter all the users' profiles who periodically posted tweets containing the keyword "CVE" from June to September 2018. By observing the temporal characteristics of tweet activities, the authors found that Twitter message streams can directly reflect cyber threats. Similar works are presented in [56,98,111–113] that learn the features of CTI from the CVE descriptions and classify each input tweet as either normal or anomalous. In particular, the framework presented in [98] aims at comparing some aspects of the information present on vulnerability databases with Twitter data. The authors start searching for tweets mentioning the vulnerabilities indexed on the vepRisk database [114] containing all entries published on NVD, PacketStorm, and other minor security databases. Moreover, they leveraged the GetOldTweets library to have access to tweets at any point in time. The pre-filtering phase is performed manually, and the final database resulted in 3,461,098 tweets. The authors of [111], instead, gather their data from Twitter by searching for tweets matching the CVE identifiers and enriching this set with additional information from the NVD database. In addition to the standard gathering of tweets about CVE IDs via the Twitter APIs, the peculiarity of [112] is that it queries the vendor sites to collect the patch release dates to determine whether a vulnerability has had immediate or deferred disclosure. Also, the paper presented in [40] uses two external sources (i.e., collections of reports, namely PRC and Hackmageddon) as a ground truth database to test the proposed method. Moreover, it leverages a large stream of tweets from GNIP's decahose, which is a real-time trend detection and discovery solution delivering a 10% sample of real-time tweets.

The paper presented in [115] deals with two social media conversation channels (Reddit and Twitter) and a collaborative software development platform (GitHub) and aims to compare user-generated content related to security vulnerabilities on these three digital platforms. In particular, Reddit is also a popular social network among researchers for its openness and its focused topic-based conversations structured around subreddits, whereas GitHub is one of the most prominent collaborative open-source software development platforms. This work shows that both Twitter and Reddit can be used to accurately predict activity on GitHub. As classically done, the dataset for this work is built by filtering posts from both Twitter and Reddit through keywords containing a vulnerability identifier (i.e., CVE) from the NVD subset. When a CVE identifier appears in a post or a comment, all the related messages (e.g., tweets, retweets, and replies) are also gathered. As for Reddit and Github, a search on subreddits (repositories, respectively) via regular expression is carried out to find the CVE IDs. Similarly, the paper [116] relies both on a rcATT database [117] and a comprehensive GitHub directory that archives APT reports dated from 2006 to 2022. To further enrich this dataset, the authors handpicked reports from 2020 to 2022 labeled with MITRE ATT&CK techniques from various cybersecurity entities such as Cyble¹ and Fortinet.² They collected a total of 724 reports (66 from Github).

Similarly to Github, *Paste sites* are Web sites or online platforms that allow users to easily share and store plain text snippets or code snippets for a temporary period. These snippets can include code, configuration

¹ <https://cyble.com>

² <https://www.fortinet.com/>

files, notes, and other text-based information. The work presented in [118], leverages data from three prevailing *paste sites* for collection based on feedback from cybersecurity experts: Pastebin, PasteFS, and Pastelink. The authors developed custom Web crawlers to collect each paste and their associated metadata (e.g., title, author).

The social Web crawler presented in [85] uses links to discussion threads on IoT vulnerabilities to traverse a forum structure and download all relevant discussions on the topic. Part of the forum employs regex-based link filters to filter out parts that are not relevant. Examples of forum crawled are: (i) Wilders Security Forums; (ii) Oracle Security Blog; and (iii) Security Forum. In the extended version of this paper, presented in [86], the author designed inTIME, a framework based on ML that can crawl data from multiple sources. As for social network monitoring, the authors show a use case where the system collects live tweets using the Twitter Stream API with keywords related to IoT vulnerabilities and leverages a CVE database as a dictionary of keywords. Finally, the authors of TI_spider [32] developed an automated data collection system from different social media, including blogs, hacking forum posts, security news, and security vendor bulletins. In particular, this system exploits 75 independent distributed crawlers, each of which monitors and collects a specific data source using a breadth-first search to collect threat descriptions. Each crawler starts the collection from a homepage, including links to threat events. For each link, it crawls the HTML source codes and extracts threat event data leveraging Xpath (XML Path language).

Table 7 summarizes, for each analyzed paper, the Social Media analyzed, how the search was initialized (which set of profiles or set of keywords are used), the type of CTI extracted, the technology used to crawl contents, the collected sample size, the type of resource gathered (i.e., tweets, users, replies, etc.), and whether the methodology is still feasible. It is observed that if crawls data from Twitter, the methodology may no longer be applicable. Since July 2023 Twitter (now called X³), has had significant changes and restrictions on access to its API, making it harder for researchers to collect and analyze data systematically. Free and open access to historical tweets is no longer available, which limits reproducibility in scientific studies.

6.3. Crawling from dark/deep web

With the term Deep Web, we refer to all Web content that is not indexed by search engines (whose content is not necessarily hidden). This includes Web pages and databases that are not accessible through search engine queries. The Dark Web, instead, represents a smaller, hidden portion of the Deep Web that is accessible through specialized software (like TOR - The Onion Router, the Invisible Internet Project, or Freenet). It typically makes use of special encryption software to hide users' identities and IP addresses and is often associated with illegal activities. The most frequently used platforms, where hackers gather to exchange malicious tools, information, and other content in the Dark/Deep Web, include hacker forums or Dark Net Forums (DNFs, for short), Dark Net Markets (DNM, hereafter), Internet Relay Chat (IRC, hereafter), and carding shops [122,123].

Hacker forums are repositories of thousands of readily available exploits enriched with comprehensive metadata, and they primarily revolve around major themes, such as carding and exploits. In contrast, DNMs often feature a substantial amount of non-cybersecurity-related content, such as pornography and illegal drugs, while lacking valuable CTI metadata, as observed by Ebrahimi et al. in their study [124]. Furthermore, researchers often face a significant risk when seeking additional details from DNMs, as products may need to be purchased. IRC and carding shops enable plain-text conversations or the posting of stolen credit card information, but they typically do not allow hackers to share exploits. Due to the analytical challenges posed by

DNMs, IRC, and carding shops, cybersecurity researchers frequently prefer to focus their efforts on hacker forums when examining exploit-related content for CTI. Consequently, a significant number of relevant studies focus on the detection of cyber threats within these hacker forums [31,86,125–127]. While hacker forums on the Dark/Deep Web are built on similar frameworks as traditional forums, they incorporate a variety of anti-crawling measures aimed at impeding large-scale, automated data collection. Common mechanisms include authentication, Turing tests, throttling, CAPTCHA images, IP address blacklists, obfuscation, paywalls, and network traffic analysis [123]. Moreover, usually, each forum framework has a unique HTML structure, naming scheme, and different internet software packages. They are accessed through Tor, a network of servers running specialized software and providing anonymity to the user. In addition, many of them are written in different languages depending on their origin (the most frequent are English, Chinese, and Russian). For all these reasons, a traditional Web crawling approach cannot be directly applied to crawling hacker forums in the Dark/Deep Web.

The general steps usually performed by all the cited articles are shown in Fig. 4, and they can be summarized as follows:

- Identifying forums. The first challenge is the identification of the relevant target forums, i.e. those containing users and content related to cybersecurity intelligence. Some researchers exploit expert knowledge to identify possible hacker forums [122,128] or leverage an already existing list of possible resources [85,126,129,130]. In many Dark Net forum conversations, participants may reference or share hyperlinks to other cyber-criminal communities or underground markets that can, in turn, be crawled [122,123] (this mechanism is referred to as snowball identification). In general, due to the underground nature of the intended targets, obtaining a curated and always updated list is quite challenging; hence, the majority of works realize custom crawlers that traverse the Dark Web to find reachable sites over Tor. For this last set, the first suitable target forums are identified by hand or by a simple keyword search to bootstrap the process [123]. After obtaining a foothold, the content of these forums is analyzed to obtain further links and addresses to other targets in a more automated fashion in later iterations [129].
- Gaining access. Since forums often require some sort of authentication to access the site, crawlers need personal accounts to log in on each site. Some sites request new users to provide only a valid email address, and others may work only if invited by other active hackers or even require users to first buy credits.
- Data collection. After the previous phases, data collection is usually fully automated. This phase deals with establishing anonymous access to the forums over Tor and the collection of raw data. Usually, a custom Web crawler is developed. The crawler will automatically download the starting seed pages, identified in the first step, and constantly discover new pages by following encountered hyperlinks. Text parser programs, using regular expressions manually identified by researchers, can be used to automatically extract meaningful information from the HTML code of the Dark Net forums (i.e., author names, thread titles, and others).

The work presented in [123] gives useful insights into the steps of Dark Net identification and Data collection. Before beginning to search for data sources, researchers must take precautionary measures to create a secure research environment to safely download, analyze, and archive Dark Net content. Researchers could have to deal with malware or malicious JavaScript code when browsing underground forums. To enhance safety, the following suggestions are given:

- Implementing virtual operating systems and networks to collect data on quarantined and isolated computers.
- Renting virtual private servers from cloud services to set up research tools on one server and then clone it to scale collection for different forums. This can provide resilience and scalability.

³ <https://x.com>

Table 7
Crawling from social media.

Paper	Year	Social media	Initial profiles/ Search keywords	Type of CTI extracted	Technology	Sample size	Feasibility
Behzadan et al. [33]	2018	Twitter	Keywords: “vulnerability”, “0day”, “DDoS”, “SQL injection”, “buffer overflow”	Tweet texts	Tweepy	21,000 tweets	✗
Sapienza et al. [104]	2018	Twitter	69 international researchers and security analysts associated with security firms	Tweet texts, author, author’s profile information, tweet location and timestamp	Twitter API	661 warnings	✗
Sauerwein et al.[111]	2018	Twitter	Regular expression in keyword search: “CVE Prefix + Year + Arbitrary Digits”	Release date, tweet text, user name, referenced websites, retweet status, date of public disclosure, vulnerability description, and CVE-ID	Twitter API	709,880 tweets	✗
Syed et al.[112]	2018	Twitter	Keyword: “CVE”	Tweet text, user follower count, user following count, user profile description	Twitter API	13,277 tweets	✗
Dionisio et al.[54]	2019	Twitter	A worldwide travel services provider, a global-company cybersecurity department, and a nation-wide power utility	Tweet text	Twitter API	5320 tweets	✗
Horawalavithana et al. [115]	2019	Twitter, Reddit, GitHub	Regular expression in keyword search: “CVE Prefix + Year + Arbitrary Digits”	Tweet text, retweets, and replies, conversation thread	Twitter API, Reddit API, Public download	105,596 tweets, retweets, replies; 170,486 posts and comments; 7240,398 activities	✗
Ji et al. [40]	2019	Twitter	Keywords provided by domain experts [119]	Tweet text	GNIP’s decahose	4975,992,550 tweets	✗
Niakanlahiji et al. [107]	2019	Twitter	Periodically automatically updated list of CTI experts who publish high-quality information about ongoing attacks	IoCs and hashtags	Twitter API	2300 tweets	✗
Rodriguez et al. [103]	2019	Twitter	Twitter accounts of hacktivists, cybersecurity feeds, security researchers, enthusiasts, and companies identified in [120] and Keyword list from the Department for Homeland Security’s “2011 Analyst’s desktop binder”	Tweet text	Twitter API	70,475 tweets	✗
Adewopo et al. [57]	2020	Twitter, Security blogs	Profiles: Brian Krebs, Cyber Security Feed, McAfee, Symantec, Hacker Combat, securityonion, CSOnline, MalwareTech, USCERT, TheHackersNews, etc.; Keywords: “ciphertext”, “cryptography”, “hacked”, “breach”, “sniffer”, “firewall”, “hijacking”, “Clickjacking”, “Malware”, “Sphearphising”, “virus”, and “vulnerability” [121]	Tweet text	Tweepy, HTML parser	500,000 tweets; 290 security blogs	✗
Alves et al. [98]	2020	Twitter	Keywords: CVE-IDs published on NVD	Tweet text	GetOldTweets	3461,098 tweets	✗
Kristiansen et al. [59]	2020	Twitter	Keywords: “malware”, “rootkit”, “cve”, “shellcode”, “sqli”, “keylogging”, “0-day”, “cyber attack”, “keylogger”, “password cracker”, “metasploit”, “reductor”, “zombie proxy”, “C&C server”, “SQL injectable”, “leaked credentials”, “remote root”, “ransomware”, “security exploit”	Tweet text	Twitter API, GetOldTweets	76,047 tweets	✗
Liu et al.[109]	2020	Twitter	Keywords provided by domain experts [119]	Tweet text and date	Twitter API	2031,766 tweets	✗
Sun et al. [97]	2020	Twitter	Regular expression in keyword search: “CVE Prefix + Year + Arbitrary Digits”	Tweet id, tweet text, CVE-ID, tweet date, screen name, coordinates	Twitter API	tweets of 2917 profiles	✗
Zhao et al. [32]	2020	Social Media	Manual search	Threat descriptions	Custom crawlers	75 blogs, security forums, etc.	✓
Chen et al.[93]	2021	Github	Manual search	Article text	Threat actions	600 articles	✓
Shin et al.[108]	2021	Twitter	35 keywords appeared with malware IOCs and 146 Twitter users (86% security experts, 12% security vendors, and 2% other security organizations)	Tweet text and retweets	Twitter API	978,414 tweets	✗
Vahedi et al. [118]	2021	Pastebin, PasteFS, Pastelink	Feedback from cybersecurity experts	Paste and associated metadata (e.g., title, author)	Custom Web crawlers	4254,453 posts	✓
Huang et al. [116]	2024	Github	Manual search	Report texts	Public download	66 reports	✓
Balasubramanian et al. [105]	2025	Twitter	Keywords: IoCs	IoCs	Tweepy	26,600 relevant tweets	✗

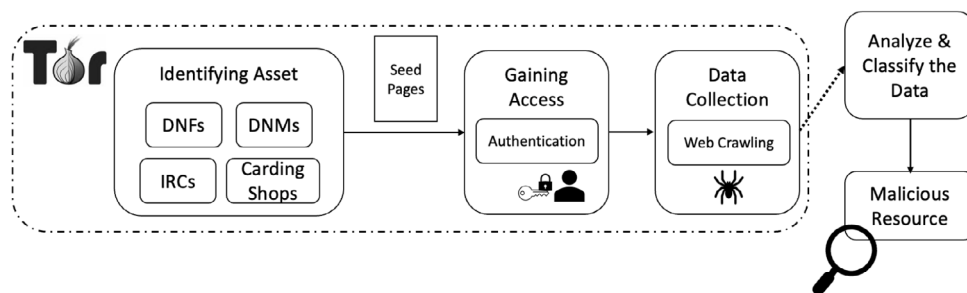


Fig. 4. General steps for crawling Dark/Deep Web.

- Check for Terms of Service (ToS) violations when identifying potential cloud service providers or other external tools.
- Sanitized Collected data can be processed for secure long-term storage.
- The database should be placed on a network distinct from the servers used for downloading Dark Net content to prevent potential malware and other security threats from compromising the integrity of the archived data.

The work of [123] presents a complete framework called DICE-E for conducting Dark Net forums identification, collection, and evaluation. Moreover, it carries out an empirical demonstration by collecting information from 4 DNFs located in the United States, China, Russia, and Iran, namely Antichat.ru, Ashiyane.org, HackHound.org, and Unpack.cn.

The paper [129] presents BlackWidow, a system that monitors Dark Web services in real-time and continuously. Moreover, it adds custom functions that emulate typing and clicking behavior to log in to forums automatically. Moreover, it employs the node.js headless Chrome browser puppeteer as a crawler within the Docker containers to collect forums' content and metadata. A similar framework is presented in [131], and it is called the Exploit Vulnerability Attention Deep Structured Semantic Model (EVA-DSSM). As for the collection part, the authors identified a large and popular exploit-specific hacker forum containing a variety of malicious tools and used a Web crawler routed through Tor to collect all exploit-category, post-date, author-name, platforms-targeted, and exploit-description data into a relational database. This resulted in 18,052 exploits across four categories (namely, Web applications, local, remote, and DoS) targeting 31 operating systems, Web applications, and programming languages.

The custom crawler called HackerRank developed by Huang et al. [132] gathers data from 5 forums (i.e., Nulled, HackThisSite, HiddenAnswers, BreachForum, Raidgets) collecting all the threads from the forum first, and then all the posts under the thread, including the username, profile, content, order, and time of the post. In addition, the authors also consider some mechanisms to deal with the anti-crawler mechanisms of the underground forums. It aims to realize an automatic method for identifying key hackers, also leveraging social network analysis metrics. In [85], the authors realize an architecture to extract CTI data not only from Clear Web and forums but also from specific Web sites on the Dark Web. This crawler is provided with several onion links that correspond to hacker forums or marketplaces, where cyber-crime tools are sold, and zero-day vulnerabilities/exploits are monitored. After a first manual authentication, the HTML text of the Web site is extracted along with useful metadata. In the extended version of this paper [86], the authors presented a complete framework capable of crawling from multiple and heterogeneous sources and supporting Dark Web crawling. This functionality relies on the use of TOR proxies to visit the user-specified onion links, and all the required actions (i.e., joining the TOR network, using the proxy, initializing the crawler) are automatically carried out via internal API calls. As for the authentication, a manual user login should be performed the first time the crawler encounters an authentication barrier. Then, session cookies are stored and used in all the subsequent crawler visits.

The works presented in [31,133] exploit data from a breached Dark Web site called Nulled.IO whose database is publicly available⁴ and from which the authors extracted hacker forum posts. Starting from this database, the authors of [31] identified the posts that are most relevant to cybersecurity through the SVM algorithm and clustered the relevant posts into topics using Latent Dirichlet Allocation (LDA). Similarly, in [134] the authors analyzed 88 million hacker forum posts of a publicly available dataset and 75,000 online articles over a 20-year timespan. In particular, they expand the CrimeBB [135] dataset,

which contains over 91 million posts collected from 32 platforms in different languages (retaining only the English one). Moreover, they include hacker forum posts collected by the AZSecure portal [136], publicly available leaked datasets [137], and data scraped from another forum, extending the CrimeBB dataset with more than 600,000 posts. The final dataset contains a total of 88,323,254 posts from 34 platforms.

The paper described in [126] performed an incremental crawling strategy bypassing the anti-crawling measures of hacker forums to collect attachments. Information was collected from 10 hacker forums, namely OpenSC, Garage4hackers, Hacksden, AntiOnline, Crackingzilla, WebCracking, SafeSkyHacks, Ashiyane, Hack, and Haker. They leveraged the internet forum software package vBulletin and allowed users to embed attachments directly into their posts that can be freely accessed. All traffic is routed through Tor, thus ensuring anonymity. After the connection to Tor is successfully established, the Python crawler begins an automatic search for attachments, performing a Depth-First Search (DFS) approach for collection. This means that the crawler starts with one subforum and crawls each topic and posts within that subforum before moving on to the next subforum.

Tavabi et al. [127] examined the dynamic patterns of discussion and identified forums with similar patterns among 80 hacker forums in the Dark and Deep Web. To do so, they exploited an already existing crawler infrastructure presented in [138]. As classically done by Dark and Deep Web crawlers, it uses anonymization protocols, such as Tor and I2P, and handles authentication to access non-indexed sites. Similarly to the previous research, the works presented in [125,139–141] leverage the same infrastructure for crawling the Dark Web and Deep Web introduced by [138]. The authors of [139] created a custom crawler and several parsers for over 200 sites relating to malicious hacking. Instead, the paper presented in [140] deals with an approach to predict the posts on hacking forums. They formulated this problem as a sequential rule-mining task, where the goal is to discover hackers' posting rules through sequences of peers' posts to make future predictions.

The authors of [125] implemented DARKMENTION, a system that retrieves information from both marketplaces where users sell information regarding vulnerabilities or exploits, and forums providing discussions about discovered vulnerabilities. They leveraged an already existing architecture introduced in [138] that includes customized crawlers and parsers built for each site and collects data from more than 400 platforms (both forums and marketplaces). They adopted ML models to ensure the collection of cybersecurity-relevant data, and the filtering phase is performed by querying the CVE IDs in the NVD database using API calls. Regular expressions are used to identify CVE mentions in the Dark Web pages. The authors of [128] realized a custom Tor-routed Web spider to crawl and download all HTML pages of a popular hacker forum. The Web spider leverages a breadth-first search strategy and a specialized Python program using regular expressions to parse all data into a local relational database. The obtained dataset with 32,766 posts (i.e., threats) made by 8,429 hackers that span a 23-year period is publicly available online.⁵ The work presented in [142] uses Dark Web data supplied by a threat intelligence company and accessed via APIs. The data is comprised of forum discussions and marketplace items offered for sale. Moreover, ML models are employed to filter out the data related to drugs, weapons, and so forth and ensure the collection of only cybersecurity-relevant data.

Differently from the papers above, which mainly gather data from hacker forums, the works presented in [57,124,143] focus on threat identification in Dark Net Marketplaces (DNMs). For all these works, the markets' identification is performed through deepdotweb.com, a Dark Net news Web site. For [143], 7 English marketplaces and 1 Russian are identified, whereas for [124], the selected marketplaces are Valhalla, Dream Market, Hansa, Alphabay, Minerva, SilkRoad3, and

⁴ <http://leakforums.net/thread-719337>

⁵ <https://github.com/HongyiZhu/>

Apple Market. A custom Web crawler is developed to traverse the onion links in a breadth-first manner. To avoid anti-crawling measures in dark net marketplaces, it implements several techniques, and it is also capable of waiting for a user's response to access CAPTCHA-protected content. Moreover, Adewopo et al. [57] collected data from two Dark Net Markets (Silkroad and Wall Street) extracted from the Arizona State University database [122]. The data set contains over 128,000 posts from different discussion threads. The thread titles are related to Carding, Newbie, Scam, Hacking, and Review threads.

The authors of [52,105,122,130] performed a crawling from multiple sources. As for [130], the different sources are: (i) 3 major DNFs (providing 204,001 threads and 14,196 authors); (ii) 5 largest DNMs (providing 224,270 product listings and 7,911 vendors); and (iii) the two largest exploit databases, such as Exploit DB and Oday.today (providing 43,678 exploit listings). The data-gathering process starts by creating an initial list of potential sites based on factors such as the number of listings, discussion threads, threat actors, and the number of listings related to cyber threats found on each site. Subsequently, several Python-based Web crawlers are designed and deployed to collect data from each of these Web sites. The authors aim to identify threats across major Dark Net data sources linking assets to threat actors, using text features and Social Network Analysis metrics. The work presented in [52] employed 11 traditional hacker forums (collected through a crawler routed via Tor), one exploit database specific for DNM (i.e., Oday.today), and several public exploit repositories collected through APIs, such as Seebug, ExploitDB, PacketStorm, Metasploit, Vulnervlab, and Zeroscience. Traditional hacker forums are crawled through a depth-first search strategy implemented for efficiency. This choice 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. The authors of [122] collected data from 51 DNFs, 12 DNMs, 13 IRC channels, and 26 carding shops. For the first step of the hacker forums identification, the authors used three approaches, namely: (i) suggestions from cybersecurity experts; they contacted both the National Cyber-Forensics Training Alliance (NCFITA), a major non-profit organization focusing on the CTI sharing and the Policing in Cyberspace (POLCYB), an internationally recognized law enforcement entity; (ii) Surface Web and Tor search engines, which are queried based on the platform names suggested by the groups of experts; (iii) snowball identification, the platforms previously identified are used as seeds for a further in-depth search. Several Web crawlers are developed to collect the raw data in HTML format for forums, DNMs, and carding shops. For IRC data, they employ two bots, emulating fake users, inside each channel. The authors of [105] built a complex crawler extracting data from both the Dark Web and Twitter. As for the Dark Web, they used multiple spiders classified based on several seed files (e.g., ACHE, Sitemap, Ahmia, Wiki, etc.). These spiders are user-defined as custom classes per the use cases' requirements and include parameters such as the URLs to be crawled, the data to be extracted, and crawling frequency. The spiders are then launched using Scrapy, an open-source Python framework to extract data using APIs.

Authors of [144] connect the Dark Net through TOR (using Polipo and Vidalia proxies) to collect information on eight large DNMs in the Dark Net, including Dream Market, Berlusconi Market, etc. The technology used to crawl and parse the Web page is the Scrapy framework. After this step, the authors designed a custom parser for each marketplace to collect important information from cybersecurity-related categories. Moreover, some studies focus on identifying key hackers in Internet Relay Chat (IRC) channels, like the paper in [145], which describes an autonomic personality analysis based on author identification in IRC conversations. Using the IRC chat logs collected through autonomic IRC bots in various cybersecurity, underground channels, and general channels (computer and politics), the authors analyze them with the Exploiting IBM Watson Personality Insights to demonstrate that the personality-based solution can work effectively in

user identification. Similarly, the work presented in [141] deals with key hackers' identification exploiting a hybrid approach that combines content, Social Network, and seniority analysis. In this work, the authors collected data provided by a commercial version of the system described in [138], from which they selected three popular English hacker forums on the Dark Web. The papers presented in [38,53] rely on an external threat intelligence platform for the Dark Web, called Sixgill, to collect hacker activity and Social Network information; then, they analyze the organizational hierarchies.

Table 8 summarizes, for each paper analyzed, the technology used to crawl content in the Dark Web or Deep Web, the identification method (if used), the type of CTI extracted, the sample size, the type of resource gathered (i.e., DNF, DNM, IRC channel, carding shops, etc.), and whether the method is still feasible.

Lessons learned. From the analysis of the works presented in this section, we can conclude that security experts and practitioners gather data from various sources such as social media, Surface Web, and Dark Web to derive Threat Intelligence. Based on the state-of-the-art research, many studies primarily focus on X/Twitter (the 80% of the analyzed works) as a key social media platform for the IoC collection. However, social media platforms are often prone to misinformation. Therefore, it is essential to validate the reliability of IoCs extracted from social media and verify the authenticity of their sources. Moreover, researchers utilize the Surface Web to collect attack incident information, such as the targeted organization, the responsible threat actor, the malware used, etc. Surface Web sources, such as vulnerability databases, can provide well-structured and reliable intelligence. Security blogs published by both vendors and independent researchers also offer information about attack incidents. However, these sources may sometimes contain biased or misleading data, necessitating cross-verification with authoritative databases. In addition, the Dark Web (e.g., hacker forums and marketplaces) serves as a critical source of operational intelligence, offering insights into threat actor communications, data leaks, and exploit discussions before public disclosure. However, accessing and analyzing Dark Web data requires specialized crawling techniques and involves significant legal and ethical concerns. Additionally, threat actors can manipulate information on these platforms, making data validation crucial. Despite these challenges, the Dark Web remains an essential resource for tracking Advanced Persistent Threat (APT) groups, monitoring cybercrime trends, and identifying potential security threats. Indeed, 47% of the analyzed works on CTI crawling include data collection from Dark/Deep Web.

7. NLP-based techniques for threat intelligence analysis

As the manual analysis of all the sources of CTI (i.e., Clear Web, Social Networks, and Dark/Deep Web) is time-consuming, inefficient, and prone to errors, several solutions to, even partially automatize the utilization of CTI information have been proposed in the literature. In this section, we aim to analyze such approaches by focusing on the NLP techniques exploited, such as text classification, text clustering, topic detection, and trend analysis.

7.1. Text classification

Text classification is one of the main tasks in the NLP field, whose objective is to categorize textual documents according to their content.

Table 8
Crawling from Dark/Deep Web.

Paper	Year	Technology	Identification method	Type of CTI extracted	Sample size	Feasibility
Almukaynizi et al. [125]	2018	Custom crawler by Nunes et al. [138]	Keyword search: "CVE Prefix + Year + Arbitrary Digits"	Article text	400 platforms	✓
Deliu et al. [31]	2018	Database online	–	Hacker forum post	Nulled.IO platform	✓
Dong et al. [144]	2018	Scrapy framework	–	Item title, description, url, vendor, vendor url, release time, shipping details, item sold, transaction details, category, CVE	8 DNMs	✓
Du et al. [122]	2018	Custom crawler, Automatic bot	Seed suggested from cybersecurity experts, surface Web and Tor search engines, snowball identification	Post text, lines of conversation, listings	51 DNFs, 12 DNMs, 13 IRC channels and 26 Carding Shops	✓
Ebrahimi et al. [143]	2018	Custom crawler	Identification based on deepdotweb.com	Product descriptions	8 DNMs	✓
Marin et al. [141]	2018	Custom crawler by Nunes et al. [138]	Manual	Topic author and content, reply author and content, topic and reply dates, user ID, user reputation	3 DNFs	✓
Marin et al. [140]	2018	Custom crawler by Nunes et al. [138]	Manual	Post text	1 DNF	✓
Shao et al. [145]	2018	Custom IRC bot	–	Author username, chat message content, time	6 IRC channels	✓
Nunes et al. [142]	2018	Data supplied by a threat intelligence company ^a	–	Posted content text (from forums), item descriptions (from markets), website name, and user name	302 Web site	✓
Tavabi et al. [139]	2018	Custom crawler by Nunes et al. [138]	manually-compiled list of 200 hacker sites	Posted message	2500,000 messages posted on a variety of Dark Web and Deep Web sites	✓
Williams et al. [126]	2018	Python custom crawler	Manual identification	User post text, attachments	10 DNFs	✓
Arnold and Samtani [130]	2019	Python custom crawlers	Initial draft of potential DNMs and DNFs manually identified	Listings, threads, threat actors, number of cyber-threat related listings	3 DNFs, 5 DNMs, 2 exploit databases	✓
Benjamin et al. [123]	2019	Custom crawler	Keyword searches; snowball collection; and deep web hidden services	Forum text content	4 DNFs	✓
Kadoguchi et al. [53]	2019	Sixgill platform	Sixgill keyword search function	Post text	3000 posts from DNFs	✓
Tavabi et al. [127]	2019	Custom crawler by Nunes et al. [138]	Manual	Post text	80 DNFs	✓
Adewopo et al. [57]	2020	Arizona State University database [122]	Thread title related to Carding, Newbie, Scam, Hacking, and Review thread	Post text	2 DNMs	✓
Ampel et al. [52]	2020	Custom crawlers	Initial manual list and snowball collection	Post text and exploit specific DNMs	11 DNFs, 1 exploit database for DNM, 6 public repositories	✓
Ebrahimi et al. [124]	2020	Custom crawler	Top markets in https://deepdotweb.com	DNM webpages text	7 DNMs	✓
Kadoguchi et al. [38]	2020	Sixgill platform	Sixgill keyword search function	Forum posts' text	1700 posts from DNFs	✓
Samtani et al. [128]	2020	Custom crawler	Cybersecurity experts	Post text, title, description, full source code, post date, target platform, attack type, and author name	1 DNF	✓
Huang et al. [132]	2021	Custom crawler	Manual	Post text, author's username, profile, content, order, and time of the post	5 DNF	✓
Samtani et al. [131]	2022	Custom crawler	Author identification of hacker forum well-known in the hacker community	Exploit text	1 DNF	✓
Suryotrisongko et al. [133]	2022	Database online	–	Post message text	Nulled.IO platform	✓
Paladini et al. [134]	2024	CrimeBB [135] dataset and other online databases	–	Post text	88,323,254 posts from 34 platforms	✓
Balasubramanian et al. [105]	2025	Scrapy framework	–	HTML Page text	1378 pages	✓

^a Cyber Reconnaissance, Inc. (CRY3CON), <https://www.cyr3con.com>.

Another popular application of text classification is sentiment analysis. Sentiment analysis tries to predict whether the sentiment expressed in a target text is either positive or negative. Text classification can be binary, as it happens for sentiment analysis or multi-class. In this case, each class is characterized by specific features present in the corpus of the text. Due to its supervised nature, text classification requires the assignment of a label to each sample in the training set, thus allowing the classifier to learn the characteristics of the different classes. To perform a classification task, any classifier needs a numerical representation of the text as input. In Table 4, we presented the most popular and effective solutions for text representation, spanning from statistical approaches, such as BoW, TF-IDF, and GloVe, to DL models, e.g., Word2Vec. The obtained representation can be used as input to traditional ML algorithms or more advanced neural networks like

Convolutional or Recurrent Neural Networks on top of a classification layer. Fig. 5 shows the basic steps of CTI text classification.

Recently, transformer-based solutions, like BERT (see Section 5.3), have gained popularity due to their impressive results. Transformer-based approaches exploit an encoder module to generate a context-accurate representation of the text, thanks to the attention system, which can, hence, be provided in input to a classification head. Text classification using the above-mentioned technologies is employed in many CTI tasks. In the following, we analyze the main applications of text classification for CTI to (i) Social Media and (ii) CTI Reports.

Social Media: In [31], the authors build a binary classifier that detects security-relevant or irrelevant posts in a hacker forum. To do so, they used an SVM classifier, achieving an accuracy of 98.82% on the task. Similarly, the authors of [33] presented a Convolutional Neural

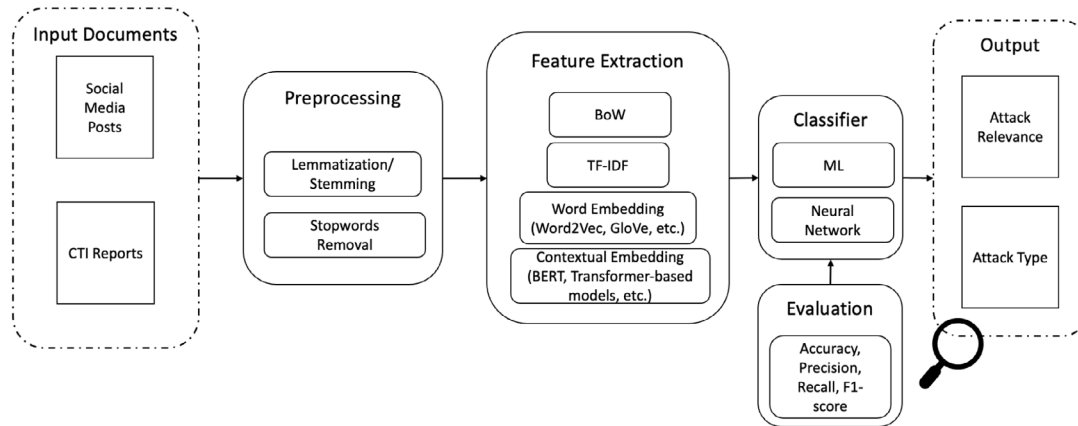


Fig. 5. General classification procedure.

Network (CNN, for short) model to perform the same binary classification task between security-relevant and irrelevant tweets. Then, if a post is classified as relevant, a second CNN is used to classify the post between 8 classes namely: vulnerability, DDoS, data leak, ransomware, 0-day, and marketing/general. The final performance of the two models is 94.72% on the binary task and 87.56% on the multi-class. In [102,146], the authors proposed a Twitter-based streaming threat monitor on a specific infrastructure that classifies whether the tweet is relevant or not for CTI. In the same way, the strategy proposed by [56] uses a TF-IDF representation technique combined with two possible ML algorithms, namely Centroid classifier and SVM, to predict the relevance of the anomaly reported in the post achieving an F1 score of the 64.3%. Hackers forums represent a rich source of data; for this reason, in [52], the authors use a C-BiLSTM neural network to perform a multi-class classification finalized to predict the label of different types of attack. In [32], the authors collected threat description data from different blogs to build a CTI domain recognizer that uses a combination of Word2Vec representation and a CNN classifier. The authors of [57] exploited data from Twitter and a hacker forum to build a logistic regression model combined with a TF-IDF representation finalized to predict the relevance of the post for security monitoring. The strategy proposed by the authors of [54] is composed of a combination of pre-trained embeddings generated by Word2Vec or GloVe and a CNN to perform a binary classification of tweets related to IT infrastructures, according to the relevance of the security-related information contained. Similarly, in [65], the authors proposed an approach to identify features, classifiers, and practices that provide the best possible detection performance for software-vulnerability-related communication in online social media channels. In particular, the strategy exploits a combination of representations obtained with Word2Vec and an SVM classifier. In [53], instead of using the representation of single words, the authors employed doc2Vec for the representation of an entire document. The results are used as input to an MLP network to extract posts that contain important information and to design proper countermeasures to the predicted attack. In [147], the authors considered informal text about ransomware from forum threads. Hence, they employed well-known algorithms like SVM, Random Forest, and Naive Bayes to detect important entities and conduct analyses. The authors of [144] developed a framework to classify cyber threats in the context of marketplaces in the darknet. To do so, after testing several strategies, including Naive Bayes and MLP classifier, the authors exploited an SVM classifier in their solution. The terminology used in forums can be used to estimate the expertise of a hacker. In this sense, the authors of [66] proposed a hacker-expertise predictor that exploits features of the text, such as TF-IDF scores of the keywords, in combination with classification algorithms, i.e., KNN and regression trees.

CTI Reports: In [39], the authors classified CTI reports according to the type of threat. To do so, they used different ML algorithms, like KNN, SVM, or decision trees, achieving performance around 80%. Another paper that analyzes CTI reports is presented in [77]. Here, the authors proposed a trigger-enhanced CTI (TriCTI) discovery system, which aims to automatically discover actionable CTI. In particular, they used a fine-tuned BERT with an elaborate design to generate triggers. The trigger vector is further trained according to the similarity with the sentence containing it. Finally, a binary prediction is performed by exploiting similarity information. Analogously, in the works [148,149], the authors employed a BERT classifier to map TTPs to the MITRE ATT&CK framework. In particular, in [148] after conducting a hyperparameter search, the authors used the MITRE dataset to fine-tune different BERT models. As an additional contribution, they compared the performance of such BERT models with a standard approach using TF-IDF and linear regression. The BERT-based models obtain an accuracy of around 80% compared to the 61% accuracy of the traditional strategy. Whereas, the proposal of Li et al. [149] consists of a label extractor component made up of two multi-label classifiers based on DistilBERT for tactics and techniques classification respectively. This system obtain a $F_{0.5}$ score of 85.50% and 75.17%. In [91], the authors compared different classifiers to predict the CTI class for an unseen cyber threat incident. In particular, they considered Naive Bayes, K-Nearest Neighbors (KNN), decision trees, and deep learning neural networks. The authors of [64] proposed a threat action extraction strategy through multi-label classification. In particular, they selected ATT&CK as a threat action taxonomy and TF-IDF as a feature. Such information is then used as input to state-of-the-art classifiers. In the same way, [62] proposed an ML model for threat report categorization. The proposed solution generalizes across reports using a combination of TF-IDF and a variety of classifiers. In [63], the authors consider the classification of bug reports similar to the descriptions of vulnerability classes from CWEs. In particular, the authors evaluated supervised and unsupervised strategies on data extracted from issue-tracking systems of two NASA missions. The proposed approach concatenates the title, the subject, and the description of each bug report and processes them through TF-IDF and a classification algorithm. The authors of [150] proposed a framework named ChainSmith, which aims at automatically extracting IoCs and their corresponding campaign stages from technical documents. To do so, they consider the most informative words according to their occurrence in the sentence, and then, these features are used to train four neural network binary classifiers that predict topic probabilities. Similarly, the work presented in [151] built a classifier that can detect posts containing IoCs and obtained an F1-score of ≈ 0.8 . A recent framework called TTPXHunter [152] aims at extracting TTPs from threat intelligence reports using the SecureBERT language model. Interestingly, this system executes a data augmentation method to build an augmented sentence-TTP dataset thus increasing the number of available samples.

7.2. Text similarity and clustering

The representation techniques described in Section 5.2, as we saw in the previous section, illustrate how pre-trained embeddings can be used as input to a classifier. An important characteristic of these representations is the intrinsic relation encoded in the vector representation of similar words. This specific aspect can be exploited to calculate a similarity score between words using distance metrics like cosine similarity. The application of this is particularly useful for dividing a set of data into clusters, giving the obtained vectors to unsupervised clustering algorithms. An example of this is described in the TOM approach presented in [58]. In particular, the authors proposed two clustering methods: (i) a semi-supervised matrix decomposition-based method, and (ii) an unsupervised term frequency-based method. The first method generates a term-malware mapping matrix from the Threat Expert reports by using the TF-IDF representation technique. The second method determines the name of the clusters based on string distances, and the malware with similar names is considered a cluster using the K-means algorithm. In the strategy presented by [153], the authors proposed an approach that improves the OSINT processing by correlating and combining IoCs coming from different feeds. In [102,146], authors performed clustering by the k-means algorithm as an aggregator of the outcome of the classifiers of tweets at the previous step of the pipeline of the approach. After the clustering, it is easier to transform unstructured data, like from social media, to standard formats, such as STIX. Similarly, in [59], the authors proposed a framework called CTI-Twitter, which applies a clustering strategy to the outcomes of the classes from a multi-class classifier. The idea is to detect inside each class the groups of semantically similar tweets to help the analysts detect trending keywords. In particular, the authors proposed two text clustering methods to group semantically similar tweets: k-means with TF-IDF and Sentence Transformer embeddings and LDA. Analogously, the authors of [81] used the representation of tweets produced by the TF-IDF in combination with the DBSCAN algorithm to cluster tweets in potential novel events. The proposed strategy uses clustering to aggregate all tweet texts in a cluster into a single corpus. Then, they performed the NER process to detect the keywords. In the proposed paper [60], the authors exploited the clustering of data from different Dark Web sources to find similar contents. The clustering process has been performed through a combination of TF-IDF and K-means algorithms to find similarities between HTML pages and detect the top-k keywords. Similarly, the authors of [109] presented a cybersecurity event discovery and evolution detection framework based on continuous tweet streams called CyberEM. In particular, the proposed framework consists of three components: pattern clustering, NMF-based event aggregation, and dynamic event inference. Focusing on the pattern recognition component, the objective is to detect the k clusters of the cybersecurity events to remove tweets that talk about cyber attacks in general. To do so, they employed *Kullback Leibler divergence*. The idea behind the approach is to detect the terms with higher KL-divergence following the institution that these terms are more informative in target domain-related tweets. In [154], the authors presented a framework called EIGER with the intent of clustering similar artifacts created by malware of the same family. The clustering strategy employs Levenshtein distance to calculate the similarity between the abstracts where the ones with the closest distance are merged sequentially. The authors of [155] formalized the threat-hunting problem from CTI reports and IoC descriptions and, in particular, used a best-effort Similarity search. The authors of [64] used TF-IDF representations to calculate semantic similarities of reports with the considered topics and actions of a threat-related article. Fig. 6 shows the basic steps of CTI text clustering.

7.3. Text mining and open information extraction

Applying Open Information Extraction (OIE, hereafter) to CTI involves using Text Mining and NLP techniques to automatically extract relevant and valuable information about cyber threats, attack

patterns, vulnerabilities, and related entities from unstructured text sources without prior knowledge or supervision.

In [156], the authors described Open-CyKG, a framework for extracting information from unstructured Advanced Persistent Threat (APT) reports and representing the retrieved data in a knowledge graph that offers efficient querying and retrieval of threat-related information. One module of this system consists of a neural attention-based OIE model to extract relation triples from unstructured APT reports. This model takes the concatenation of all inputs and passes it to two Bidirectional Gated Recurrent Units (Bi-GRU) layers, followed by an attention layer, two Time Distributed Dense (TDD) layers, and finally, a SoftMax layer for prediction. The work proposed in [99] develops a system leveraging both a self-attention deep neural network (SA-DNN) model and a text mining approach to identify the vulnerability category from the description text contained within a report. Similarly, the proposal of [29] aims at extracting features from unstructured CTI reports and attributing them to cyber-threat actors. Detailed feature sets, i.e., TTP, tools, malware, target organization, country, and application, have been used. Moreover, in this work, the authors proposed a novel embedding model called attack2vec to extract features from unstructured CTI reports. This novel model achieves accuracy, precision, recall, and F1-score of 96%, 96.5%, 95.58%, and 95.75%, respectively.

The paper presented in [157] deals with the design of software for the automatic extraction of actionable threat intelligence from raw log data. The authors employed anomaly detection to disclose unknown attacks, and they pursued the iterative clustering and enrichment of anomalies with optional human verification to transform low-level log events into complex attack patterns. In this framework, every anomaly is primarily defined by the feature that triggered the detection mechanism. Then, the anomaly is transformed into an alert with enriched attributes storing contextual data as well as event and execution information. Finally, they eventually generate clusters of alarms that frequently occur together.

Finally, the proposal of Chen et al. [158] describes a method for identifying the security-sensitive levels of an API call by matching the run-time parameters with an IoC database extracted from CTIs. To do so, the authors leverage an automatic IoC mining that analyzes textual CTIs and includes three steps, i.e., potential IoC filtering, IoC recognition, and IoC extraction. The IoC mining step of the system includes a classifier that classifies each potential element as IoC or not.

7.4. Cross-lingual threat intelligence

Cyber threat actors operate on a global scale, and many of them communicate in languages other than English. To comprehensively understand and combat cyber threats, it is essential to monitor and analyze threats originating from communities with different languages. However, while text classification techniques have been mainly used for cyber threat detection in English platforms, this task is hindered in non-English ones due to the language barrier and lack of ground-truth data. Recent approaches usually filter out non-English data or make an automatic English translation to leverage monolingual models and overcome this issue. For instance, in all the works based on a Twitter crawler, the authors usually start the search with an English keyword or set a condition on the preferred language of the post or its country of origin (i.e., “lang: en”). In a popular work of Samtani et al. [159], Russian data was machine-translated to English using Google Translate. However, translation errors can deteriorate the classification results, and on the other hand, training separate monolingual models on low-resource non-English languages is impracticable.

For the above reasons, the authors of [160] demonstrated that their deep cross-lingual model can jointly learn the common language representation from two languages (namely, English and Russian), and it outperforms a monolingual model trained through machine-translated data for identifying cyber threats in non-English DNMs. They randomly sampled the product descriptions in each language by preserving the

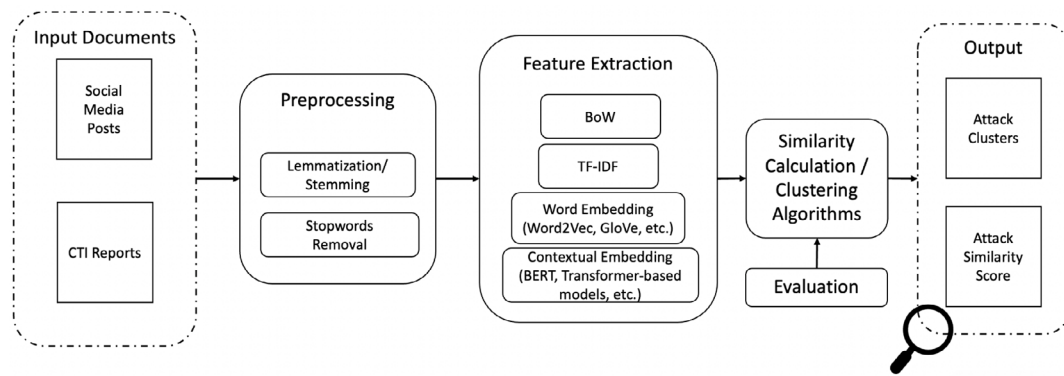


Fig. 6. General clustering procedure.

ratio of cyber to non-cyber products, and then the manual labeling step as cyber threats or non-threats was performed by cybersecurity experts and a Russian speaker. Moreover, the CL-LSTM (CrossLingual LSTM) algorithm leverages BiLSTM to jointly learn the common hacker language representation from English and Russian DNMs. The authors of [161] collected data from an English forum, a Russian forum, and two French forums and designed A-CLKT, an adversarial learning procedure that learns language invariant representations across two languages automatically. In the first phase of automated text representation, A-CLKT reserves an LSTM for each language to automatically create an embedding of hacker forum text. In the second phase, the authors leveraged a GAN to devise a novel adversarial learning strategy to operate on the English and non-English representations. Finally, a classification step is performed to classify the text as a threat or non-threat.

The work presented in [162] collects data from Twitter both in English and in Russian, exploiting Twitter API language capabilities to detect language through the related flag (en=English, ru=Russian). Through two lexical databases, Wordnet and Russnet, the authors create relationships between the two languages' cybersecurity words and align them in vector embeddings. Then, an LSTM-based neural machine translation architecture is used to translate cybersecurity text from Russian to English. Finally, the work by Wu et al. [95] adopted a Language Technology Platform (LTP [163]) that is an integrated Chinese language processing platform that includes a suite of high-performance NLP modules and relevant Chinese corpora.

7.5. Topic detection and trend analysis

Topic modeling is an unsupervised ML technique used to automatically identify topics in a collection of text documents and discover patterns in groups of words or documents that are semantically similar. Two of the most commonly used algorithms for topic modeling are Non-negative Matrix Factorization (NMF) and LDA [164]. NMF is a statistical method that aims to reduce the dimension of the input corpora, using the factor analysis method to provide comparatively less importance to the words with less coherence. LDA, instead, computes the posterior probability of the distributions of topics per document and the distribution of words per topic. These estimations are presented as matrices and used to infer two outputs, (i) the dominant topic per document and (ii) an ordered list of words that constitute that topic. The works that base their topic modeling approach in LDA are the following [31,99]. Specifically, the authors of [31] applied LDA after using SVM to remove irrelevant posts to extract ten topics. The interpretation of each topic's meaning is based on reviewing its top documents(posts) by human operators. The extracted topics deal with leaked credentials, malicious proxies, undetected malware, and asset-specific CTI. The work by Vishnu et al. [99] performs topic modeling on a 13-class categorization problem using the LDA method to discover

the topic trends within the dataset for the various CVE categories. In particular, summarized vulnerability descriptions are converted to a document-term matrix and given to the LDA model to create the topic model. The model produced ten topics, from which the authors manually selected the most suitable to describe each vulnerability category.

As seen before, the topic generated by LDA may not be meaningful to a user. Hence, the authors of [165] apply to CTI the algorithm SeededLDA proposed by Jagarlamudi et al. [166], which allows a user to give additional information to the topic model to learn topics of specific interest to a user. In SeededLDA, a user provides seed sets according to the state and the environment of the organization (i.e., IoT or Financial Industry) to guide the topics together with security blog posts.

The paper by Hossen et al. [167] utilizes both LDA and another popular algorithm for topic modeling: Non-negative Matrix Factorization (NMF). They collect and analyze data from a well-known hacker forum to identify and classify possible CTI. The paper described in [33] leverages topic modeling as the preliminary step of manual annotation of tweets to speed up and increase the accuracy of the whole process. The topic modeling API of IBM's Watson Natural Language Understanding service (WATSON) [110] is used to recollect text classification into five categories connected to cybersecurity for the textual contents of each tweet. The text category assignment was restricted to the top three categories with the highest confidence score.

The approach proposed in [61] describes an improved keyword feature extraction method, namely the ITFIDF-LP (Incremental TF-IDF method considering word location and part of speech) method. In particular, this strategy considers keyword identification tools, whether a word is a stop word, its position in the article, and which part of speech this word represents.

Given the limitations of prevailing topic models, several works developed custom approaches based on LDA [118,133]. More recent topic modeling approaches (namely, BERTopic and Top2Vec) are leveraged by [133]. BERTopic relies on pre-trained transformer-based language models to compute document embedding, cluster these embeddings, and generate topics with the class-based TF-IDF (Term Frequency-Inverse Document Frequency) procedure [168]. Top2Vec is a new topic modeling and semantic search algorithm that creates embedded document/word vectors, makes lower dimensional embedding, finds dense areas of documents, calculates the centroid, and finally finds the closest word vectors [169].

To categorize long and contiguous text (e.g., pastes), the authors of [118] incorporate BERT into LDA, proposing the BERT-LDA framework. It consists of three main components, namely (i) BERT's encoder that tokenizes each word in each sequence (sentence) for every input paste; (ii) BoL model (produced by BERT) replacing the traditional BoW in the conventional LDA model captures information at about each paste's semantics at the sequence-level, rather than at the word-level; and (iii) Topic Generation using LDA produces topics based on

Table 9
Tool/Technique for topic modeling.

Paper	Year	Tool/Technique
Behzadan et al. [33]	2018	WATSON service [110]
Deliu et al. [31]	2018	LDA
Nagai et al. [165]	2019	SeededLDA
Li et al. [61]	2020	ITFIDF-LP
Hossen et al. [167]	2021	LDA and NMF
Sleemans et al. [170]	2021	DTM
Vahedi et al. [118]	2021	BERT-LDA
Suryotrisongko et al. [133]	2022	BERTopic and Top2Vec
Vishnu et al. [99]	2022	LDA
Huang et al. [116]	2024	BERT

each paste's BoL. The proposed BERT-LDA model was applied to all pastes gathered from Pastebin, PasteFS, and Pastelink platforms. The authors manually assigned names to five prevailing topics extracted by the model (namely, hackers, malware, networks, Web sites, and PII) and checked the results by comparing them with the selected keywords. Similarly, the work of Huang et al. [116] detects MITRE ATT&CK techniques and TTP in CTI reports. To do so, in the first stage this framework performs the training of a topic classifier to sift through sentences in a CTI report, identifying those that specifically relate to cybersecurity techniques by fine-tuning the BERT model.

The authors of [170] used the Dynamic Topic Model (DTM, hereafter) to model multiple document collections over time. They exploit Wikipedia concepts related to cybersecurity as a context model for training the DTM to understand how the concepts found among documents are changing over time. The leveraged corpus belongs to arXiv Cryptography and Security research papers. An important contribution of this work is the automatic domain concept extraction using Wikipedia concepts. They captured 3,836 concepts from Wikipedia, exploiting them to establish the context for the topic modeling portion.

Table 9 summarizes, for each analyzed paper, the tool and/or the technique employed by the authors to perform topic modeling.

7.6. Text summarization

Text Summarization is often used to reduce the verbosity of a document or Web content and obtain a concise description of the attack behavior that can be directly used to detect the attack itself. Indeed, threat reports are usually characterized by a significant amount of complex and irrelevant text, and only a small portion of the report describes attack behavior. Although topic classification has been used to identify topic-related context among out-of-domain contexts (e.g., advertisement text versus technical text), they are not good enough to describe observable attack actions discriminating among technical concepts. Moreover, inside each sentence, some parts of the speech, such as adverbial and adjectival, do not contribute to the behavior description of the attack, and that can be safely removed. To solve this issue and provide a good summary of CTI content, authors of [83] designed a two-step approach that consists of a BERT classifier, which deals with sentence verbosity, and a BiLSTM network, which deals with word verbosity. In particular, for the first step of sentence verbosity, they use a 12-hidden layer BERT to discern the productive sentences from the non-productive ones. To train the model, they used 8,000 labeled sentences. For removing the word verbosity, instead, understanding the words' roles in the text summary is crucial, and for this reason, they used a re-implementation of a deep BiLSTM model [171]. Since this model was not fine-tuned to handle cybersecurity sentences, they trained the model using 3,000 manually labeled sentences. Similarly, the paper presented in [92] performed text summarization, adopting the BERT transformer model to perform word embedding and BiLSTM to extract threat entities.

The authors of [172] proposed a lightweight scheme, called CVErizer, for summarizing CVE's content, which consists of two steps: (i)

information extraction and (ii) classification of vulnerabilities. They observed that the most relevant information pieces usually contained in a CVE description are: the vulnerability name, the name and the version of the software affected by the vulnerability, the kind of attacker who could exploit the vulnerability, the mechanism used, and the effect. This recurrent pattern is then formalized in an XML file, and through a lightweight taxonomy, these vulnerabilities are categorized into different types, reaching an accuracy score of 81%. To monitor the cybersecurity events, Li et al. [89] summarized CTI text using *TextRank*, a graph-based ranking algorithm. This approach entails the transformation of natural language text into a graph, where individual sentences are considered as vertices, and connections are established based on sentence similarity. The similarity is determined by calculating the count of common tokens found in the lexical representations of two sentences and then dividing this count by a normalization factor. Sentences exhibiting greater connectivity to other vertices are accorded higher importance and are subsequently ranked, resulting in the extraction of pivotal sentences. Arranging these summaries chronologically using timestamps reveals the development of cybersecurity events.

Lessons learned: From the studies presented in this section, we can conclude that data collection from open-source platforms may lead to the inclusion of irrelevant information. Text classification plays a crucial role in generating relevant CTI by filtering out such irrelevant data. Dark Web platforms, a primary source for attack tools and malware, often use non-English languages for communication. Some studies have explored cross-lingual threat intelligence to overcome language barriers

8. NLP-based techniques for relation extraction

In this section, we delve into research dedicated to extracting relationships from cybersecurity data. To perform relationship extraction, multiple researchers have pre-identified cybersecurity entities. Furthermore, certain studies have successfully identified cybersecurity events from these recognized entities and extracted valuable relationships. Additionally, we have reviewed various research works that center on the development of graphs using this extracted knowledge. The process of relationship extraction and constructing a Knowledge Graph (KG) from security text can significantly contribute to generating actionable CTI.

8.1. Named entity recognition

Named Entity Recognition (NER) assists in recognizing entities (proper nouns) from the unstructured text and serves as a fundamental technique for tasks such as information retrieval, constructing knowledge bases, question answering, generating automatic text summaries, and providing semantic annotations. The named entities include real-world objects like persons, organizations, locations, dates, and other entities with proper names. NER holds a pivotal role in NLP by focusing on identifying and categorizing named entities within the text.

There are general tools available, such as Spacy, Stanford NER, NLTK, Flair, etc., that can recognize generic entity types like locations, persons, organizations, and dates. However, these entities might not fulfill the specific demands of specialized domains like cybersecurity [173]. Cybersecurity requires the identification and classification of entities unique to the field, including malware types, operating systems, attack methods, and other terms specific to cybersecurity. Fig. 7 displays a sample of cybersecurity text sourced from the MITRE Web site,⁶ along with its corresponding entities. The significance of identifying and labeling these domain-specific entities becomes evident in various cybersecurity tasks. For instance, in the context of malware analysis, NER can pinpoint the types of malware within a text

⁶ <https://attack.mitre.org/software/S0268/>

Bisonal TOOL has been used by Tonto Team THREAT GROUP against public and private sector organizations in Russia LOC, South Korea LOC, and Japan LOC since at least Dec 2010 DATE.

Fig. 7. Cybersecurity named entity.

document [174], thus aiding analysts in comprehending the threat's nature. In the classification of attacks and vulnerabilities, NER plays a role in identifying the precise types of attacks or vulnerabilities [82] mentioned in security reports or incident logs. Additionally, NER contributes to the creation of cybersecurity KG and structures that organize and present cybersecurity-related information systematically.

8.1.1. Annotation strategy

Recently, NER has been predominantly executed through supervised learning classifiers that require labeled text as input. Since the rapidly evolving threat landscape is characterized by emerging attack mechanisms and the convergence of various attack and malware types, the imperative arises to develop customized NER datasets. In the NER task, each word of unstructured text is assigned a specific tag instead of assigning a label to an entire paragraph. Researchers have employed various tagging schemes and annotation tools for annotating this task.

Tagging Schemes: Researchers explored tagging schemes, including BIO, IO, IOE, BIOES, BILOU, etc. Among these schemes, BIO and BIOES are extensively utilized in CTI.

- **BIO or IOB:** BIO stands for “Begin, Inside, Outside,” also referred to as the IOB format. The B (Beginning) and I (Inside) tags are primarily employed to label relevant entities. The B index indicates the start of a specific entity instance, while the I index signifies the continuation of the particular entity. Additionally, an O tag is used to denote terms that exist outside of any entity context. For instance, consider the sentence extracted from the Bleeping Computer site: Major US energy organization targeted in QR code Phishing attack. In this example, US is tagged as B-Location, Energy is assigned with the B-Industry tag, QR is labeled as B-attack type, code as I-attack type, Phishing as I-attack type, and all other tokens are assigned the O tag. In studies [55,72,75,96,156,175–180], researchers employed the BIO/IOB tagging scheme for annotating the tokens as cybersecurity entities.
- **BIOES:** The BIOES scheme also employs BIO tags akin to the BIO approach but contains two additional tags: E and S. The E tag indicates the end of a specific entity, while the S tag is utilized for single-word entities. In the sentence, Major US energy organization targeted in QR code Phishing attack, US is designated as S-Location, Energy is assigned the S-Industry tag, QR is marked as B-attack type, code as I-attack type, Phishing as E-attack type, and all other tokens are labeled with the O-tag. Few researchers have adopted the BIOES tagging scheme in their studies [46,181,182].
- **BIOFR:** Traditional BIO tagging typically concentrates solely on identifying entities and does not consider the sentence's structure. In contrast, BIOFR [183], which stands for Begin (B), Inside (I), Outside (O), Front (F), and Rear (R), incorporates attention to words surrounding the entities. However, it is worth noting that BIOFR does not regard nouns as responsible for threats or vulnerabilities; instead, it summarizes them using verbs. Take the sentence, “A stack out-of-bound error was flagged in ...” When we apply BIOER, we get tagged sentences like this: “A [F-B-Tri] stack [B-Tri] out-of-bound [I-Tri] error [I-Tri] was [R-B-Tri] flagged [R-I-Tri] ...”

Annotation Tools: For expediting the manual annotation process, the researchers employed various annotation tools, including *Prodigy*, *BRAT*,

YEDDA, *Doccano*, *LightTag*, etc. The BRAT Rapid Annotation Tool (BRAT) [184] serves as a Web-based text annotation mechanism primarily designed for NLP tasks, encompassing both entity and relation annotation. The authors opted for BRAT in their work [75,175,180,181]. Prodigy [185], a commercial tool, supports not only text-based tasks like NER, dependency parsing, and part-of-speech tagging but also tasks for audio and video annotation and computer vision. Few works [173,186] used Prodigy for annotating the text. YEDDA [187], on the other hand, is an open-source lightweight tool supporting command-line annotation. It supports languages such as English and Chinese and works with symbols and even emojis. YEDDA is adopted by Ren et al. [178] in their study. Further, Zhou et al. [72] employed YEDDA for annotating the Chinese security-related text.

8.1.2. NER approaches

Conventional Methods: Early research on NER relied on rule-based and statistical machine-learning models. Rule-based strategies hinged on rules crafted by domain experts, amalgamating gazetteers (curated lists of known entities), and syntactic-lexical patterns [188]. Hanks et al. [173] employed gazetteers to identify a range of entities, including protocols, malware types, programming languages, attack types, file extensions, and operating systems. These gazetteers were constructed by querying data from Wikidata. Additionally, the authors employed regular expressions to detect indicators such as CVEs, ports, hashes, and IP addresses. For recognizing URLs and Email addresses, they utilized the Spacy tool. In [55], the authors adopted hybrid techniques, incorporating rule-based and DL-based methodologies to identify cyber entities effectively. Specifically, they extracted CVE IDs, commonly employed to denote vulnerabilities within software or products, using regular expressions and replacing them with the CVE tag. Further, in [75], Alam et al. adopted regular expressions to recognize indicators such as IPv4, CVE, SHA1, SHA256, Email, and File Path. On the other hand, statistical methods harnessed the power of ML algorithms such as Perceptron, Support Vector Machine (SVM), Hidden Markov Model (HMM), and Conditional Random Field (CRF). The most effective methods for NER often involve the utilization of CRFs. CRFs address the NER task by implementing a sequence labeling model, where the assignment of a label to an entity is intricately linked to the labels assigned to neighboring entities within a predetermined contextual window. However, nowadays, researchers are using DL classifiers instead of ML classifiers.

Deep Learning for NER: The traditional approaches involved meticulous manual intervention in creating rules and undertaking intricate feature engineering. Following the rise of deep learning, many researchers have embraced DL-driven techniques in their research on NER to circumvent the laborious task of feature engineering. Significantly, the transformer framework showcases remarkable efficacy across diverse NLP tasks with models like BERT. In [54,189], authors explored BiLSTM with CRF for identifying cybersecurity entities. In the study [190], the authors employed LSTM recurrent neural networks, a domain-agnostic approach, in combination with a CRF method to extract entities within the cybersecurity domain. VIEM [82] used a NER mechanism to identify vulnerable software names and versions. The NER model developed by Bidirectional Recurrent Neural Network achieves 97.8% and 99.1% precision and recall, respectively. Furthermore, they employed a heuristic approach that incorporated dictionary lookups to rectify inaccuracies in the NER model. With the introduction of a gazetteer, the overall accuracy increased to 99.69%. In [191], Georgescu et al. developed an automated system aimed at serving as a semantic indexing solution to identify pre-existing vulnerabilities within IoT environments. They created a domain ontology and crafted a NER model to conduct semantic text analysis. To train the NER model, they employed the Watson Knowledge Studio tool, which leverages a semi-supervised learning approach. Pingchuan et al. [189] generated word embedding vectors by simultaneously training the model and the word vectors with the help of one hot encoding. The study by

Kim et al. [55] integrated character-level attributes through Bag-Of-Character representations and incorporated GloVe for word embedding. Their strategy encompassed using a BiLSTM model along with CRF to extract cybersecurity entities. Similarly, Gao et al. [179] adopted a BiLSTM-CRF architecture, which incorporates domain dictionary embedding and a multi-attention mechanism for enhanced performance. They constructed a list of 15,357 words for domain embedding by collecting information from security blogs, CVE, NVD, and Wikipedia. In [192], the authors employed a novel approach to identify entities in both English and Chinese text. This method leveraged a neural network, CNN BiLSTM CRF model, in conjunction with a Feature Template (FT). The FT serves to extract local contextual features, while CNN was employed to capture character-level feature sequences, and BiLSTM was used to create the global feature vector.

The study in [45] assessed various DL-based NER algorithms across a wide spectrum of sentence types, encompassing concise and lengthy structures. They investigated both domain-independent Word2Vec embeddings and domain-specific Word2Vec embeddings (trained on a cybersecurity corpus). Their findings demonstrated that the utilization of domain-specific embeddings resulted in superior performance. Further, they pinpointed that the optimal choice for cyber NERs entails the fusion of BERT embeddings with Bidirectional LSTMs. In [181], Wang et al. introduced APTNER, a NER dataset comprising 21 entities derived from STIX 2.1. The authors evaluated multiple NER models and discovered that the BERT embedding-based BiLSTM model with CRF achieved the highest F1 score of 82.31%. Also, in [96], the authors employed a combination of BERT, BiLSTM, and CRF for entity extraction and achieved an F1 score of 0.949. In another study, Wang et al. [175] leveraged knowledge engineering to boost the performance of their NER model, which relied on a combination of BERT-embedded BiLSTM and CRF. The model built using knowledge engineering achieved an F1-score of 91.62%. Wang et al. [46] developed a NER model using a novel Neural Network Cell called GARU, which combines a Graph Neural Network and a Recurrent Neural Network. They concatenated diverse embedding dimensions such as character level, word level, POS, and dependency relation in the embedding layer. For word embedding, they used two methods: Word2Vec, trained on cybersecurity datasets [193], and PERT [194], a BERT-inspired auto-encoding model. They incorporated character-level features from character-based CNNs. Additionally, they created a module for detecting entity boundaries and predicting entity heads and tails. Sun et al. [84] employed supervised neural network approaches to extract and categorize 15 cybersecurity entities. Their approach commenced with the use of Stanford CoreNLP to segment sentences in articles based on punctuation marks. Subsequently, token segmentation was performed within those sentences. They combined word embeddings with feature embeddings and inputted them into a BiLSTM RNN. They conducted an evaluation and performance comparison of their entity recognition method with two approaches: one using LSTM with ELMo embeddings and the other with fine-tuned BERT embeddings. The neural network incorporating ELMO and BERT embeddings achieved exceptionally high accuracy levels, reaching 99.8% and 99.2%, respectively. In [41], various fine-tuned models, including BERT, RoBERTa, ELECTRA, and DeBERTa, were utilized to perform NER and extract IoCs. Token annotation followed the CoNLL-2003 structure, and IoCs were extracted from sentences based on predefined rules for 16 entities using the `pyarsing` module. When evaluating model performance, the fine-tuned DeBERTa model demonstrated superior results compared to the other models. However, in practical case studies, it was observed that the ELECTRA model effectively identified new IoCs. While Mouiche et al. [195] experimented with a SecureBERT classifier combined with a BiLSTM layer and a Time-Distributed Dense (TDD) layer to accurately identify threat entities within reports. They trained their model using the publicly available DNRTI dataset and validated it against DNRTI, CyNER, and STUCCO. Their approach achieved F1-scores of 0.90 on DNRTI [181], 0.96 on CyNER [75], and 0.98 on STUCCO [196].

The enhancement of the identification and categorization of cybersecurity-related terms is the aim of Liu et al. who proposed a semantic augmentation approach in [197]. Their approach consists of three primary elements: internal domain augmentation, external general augmentation, and mixed linguistic features. Internal domain augmentation enriches the meaning of input words by incorporating insights from the K most similar words within the cybersecurity domain. Additionally, they fine-tuned the BERT model and utilized the resultant token representations for external general augmentation. Moreover, they generated a composite set of linguistic features, encompassing elements such as POS, morphology, and other component-related features. In [182], Yang et al. underlined the significance of improving word information in the context of Chinese cybersecurity NER. They have introduced a novel model called Star-HGCN, in which the hybrid embedding layer enriches word information by integrating character-level data and POS embeddings. In this model, the Star-Transformer layer adeptly captures implicit semantic connections among words in a sentence, while the GCN layer is applied to model explicit syntactic dependencies, drawing from dependency tree structures. In the decoding phase, a CRF is employed to carry out sequence labeling for sentences.

Alam et al. [75] developed an open-source toolkit named CyNER to extract cybersecurity-related entities from unstructured text. To accomplish this, they combined various methods: transformer-based models for identifying cyber entities, heuristic techniques like regular expressions for indicators, and established models like Spacy and Flair for general entities such as locations and persons. They gathered 60 threat reports from MITRE ATT&CK to generate a benchmark dataset. Using the BRAT annotation tool, they labeled the dataset with different types of entities, including vulnerabilities, organizations, systems, indicators, and malware. The transformer-based model, specifically XLM using RoBERTa large, achieved an F1 score of 76.66% on this task. In [105], the authors utilized CyNER data for training, primarily developing the NER model using BERT to extract IoCs. Most research in NER relies on labeled data, but in cybersecurity, such data is scarce. So, In [198], the authors introduced an approach that combines a Generative Adversarial Network (GAN) with the BiLSTM Attention CRF model to acquire labeled data from crowd annotations. The GAN is utilized to identify shared features within crowd annotations, which are then integrated, with domain dictionary features and sentence dependency features, into the BiLSTM-Attention-CRF model. This integration is aimed at enhancing the quality of crowdsourced annotations. Similarly, In [199], the authors proposed an adversarial active learning approach for cybersecurity NER. They employed a BiLSTM layer to encode word embeddings and then utilized an additional LSTM layer to decode hidden representations obtained from the dynamic attention layer. More specific information about the representative work on neural network-based NER is included in Table 10. The table lists the different data sources the researchers employed as well as the embedding technique they applied.

8.1.3. NER corpus

Generating actionable threat intelligence through DL-based NER necessitates labeled data. Based on different tagging schemes and annotation tools, researchers invest significant effort in annotating security reports to create a NER corpus. Researchers defined different entity types according to their specific threat intelligence objectives. Few research works explain how they conducted and validated the annotation. In the study by Kim et al. [55], the annotation was carried out by five cybersecurity experts who were assigned text segments ranging from 1500 to 5000 lines. In another study [181], the authors enlisted the assistance of 30 undergraduates and four graduate students to label security-related text and verify the labeling process. Another study [175] involved three graduate students for cross-verification of the annotated data.

To evaluate the quality of the annotation, Hanks et al. [173] conducted an inter-annotator agreement assessment. Notably, during the

Table 10
Research work on Named Entity Recognition.

Paper	Year	Data source	Entities	Tagging scheme	NLP technique	Result	Remarks
Gasmi et al. [190]	2018	Bridges et al. dataset [196]	vendor, application, version, file, OS, hardware, edition	BIO	Word2vec	LSTM + CRF - 83.37% (F1 score)	Poor performance on the edition and hardware
Dionísio et al. [54]	2019	Twitter	Organization, product or asset, version number, vulnerability, ID	Entity label+ O tag	GloVe, Word2Vec	BiLSTM - 92% (Average F1 score)	Number of entity type is less
Dong et al. [82]	2019	CVE, vulnerability report	vulnerable software name and version	Entity label+ O tag	Skip-gram	Bidirectional RNN - 0.978 (precision) and 0.991 (recall)	Extracted vulnerable software name and version
Qin et al. [192]	2019	WooYun vulnerability database and Freebuf Web site	VUL ID, network relevant term, software, organization, location, person	BIO	Word2Vec	FT+CNN+BiLSTM+CRF - 86% (F1 score)	Detection of entities in both English and Chinese languages
Dasgupta et al. [45]	2020	Microsoft security bulletin, Adobe security updates, NVD Twitter, Technical reports, CVE	STIX 2.0 + Exploit-Target	NM	Word2Vec, BERT	BERT+BiLSTM+CRF 88.60% (F1 score)	Trained on dataset with varying text lengths and complexity.
Kim et al. [55]	2020	PDF documents	Main entities: Malware, IP, Hash, Domain/URL, Total 20 sub-entities	BIO	Bag-Of-Character	BOC+BiLSTM+CRF 75.05(F1 score)	CVE achieves 100% performance and .url.normal, url.unknown, malware.unknown, exhibit high differences in precision and recall.
Pingchuan et al. [189]	2020	Lal et al. dataset [200]	Software, modifier, OS, consequences, attack, means, file name, network, hardware	BIO	NM	BiLSTM + CRF - 89.38%(F1 score)	Low f1 score for hardware and network types
Gao et al. [179]	2021	NER: Bridges et al. dataset [196], Domain dictionary: Wikipedia, CVE& NVD, blog	Application, version, hardware, OS, file, vendor, edition	NM	Domain dictionary embedding(N-gram feature)	BiLSTM+Dict+Att+CRF 88.36%(F1 score)	Lowest performance on hardware and edition
Li et al. [199]	2021	APT report, CVE, security news, blogs	organization, location, software, malware, indicator, vulnerability, course-of-action, tool, attack-pattern, campaign	NM	Word2Vec	Dynamic-att+ BiLSTM+ LSTM - 45% sample achieves 88.27% (F1 score)	Adversarial learning to reduce labeling effort
Alam et al. [75]	2022	MITRE ATT&CK	Malware, indicator, system, organization, vulnerability, email, URL, file path, hash, CVE, generic entity types of spacy and flair	BIO	BERT	XLm+RoBERTa - Large 76.66%(F1 score)	Developed CyNER python library
Chan et al. [41]	2022	AlienVault OTX	attack technique, bitcoin address, CVE, Microsoft defender threat, domain, email, MD5, SHA-1, SHA-256, file path, hostname, IPV4, IPV6, sslcert fingerprint, URI and URL	BIO	BERT	DeBERTa-97.51% (F1 score)	The fine-tuned ELECTRA model demonstrated the ability to extract novel IoCs accurately.
Sun et al. [84]	2022	Security news article, vulnerability database, Twitter	organization, person, device, product, Web site, capability, file, malware, money, number, purpose geopolitical entity (GPE), time, CVE, vulnerability, version	BIO	POS, ELMo, BERT	ELMO + BiLSTM RNN - 99.8% (F1 score)	ELMO embedding better than BERT
Wang et al. [181]	2022	APT reports	APT, SECTEAM, IDTY, OS, EMAIL, LOC, TIME, IP, DOM, URL, PROT, FILE, TOOL, MD5, SHA1, SHA2, MAL, ENCR, VULNAME, VULID, ACTION	BIOES	ELMo, BERT	BERT+BiLSTM+CRF 82.31% (F1 score)	Generated a publicly available NER dataset
Wang et al. [175]	2022	Twitter, Blogs, Forums, Cybersecurity company sites	Hacker groups, sample files, malicious samples, security teams, attack time, tool, country, industry, organization, user, methods, vulnerability, mode of transmission, file type, alias	BIO	Word2Vec, GloVe, ELMO, GPT, BERT	BERT+ BiLSTM+ CRF+ Knowledge engineering 91.62% (F1 score)	Performance of vulnerability recognition increased by the Knowledge Engineering
Yang et al. [182]	2023	FreeBuf portal and blogs	Vulnerability ID, Computer terminology, Software, Organization, Location	BIOES	Word, Character level, POS embedding	Star-HGCN -90.98% (F1 score)	Performance of organization and software improved by Star-HCN method
Wang et al. [46]	2023	Bridges et al. dataset [196], OpenCyber: APT reports, CVE, security bulletins	OpenCyber: Threat actor, Attack pattern, Vulnerability, Software, Malware, Campaign, Indicator, Course of action, Tool	BIOES	Word2Vec, PERT	PERT based FIEBD: Bridges et al.- 94.56% OpenCyber - 87.34%(F1 score)	Predicted entity boundaries(start and end points)
Arikkat et al. [96]	2024	Security blogs	Asset, Attack, Attack pattern, Campaign, Date-time, Identity, Indicator, Infrastructure, Location, Malware, Threat actor, Tool, Vulnerability	BIO	BERT	BERT+BiLSTM+CRF - 94.9% (F1 score)	Entity types are aligned with STIX entities
Mouiche et al. [195]	2025	DNRTI [181]	HackOrg, OffAct, SamFile, SecTeam, Tool, Time, Purp, Area, Idus, Org, Way, Exp, Features	BIO	BERT	SecureBERT+BiLSTM+TDD-Att - 90% (F1 score)	Entity extraction performed using SecureBERT

* NM indicates Not Mentioned.

initial round of annotation, the authors found that less than 50% of the annotations reached a consensus. The primary source of disagreement

stemmed from varying interpretations of entities, such as distinguishing between software names and tools or products. For instance, when

Table 11
Publicly available NER corpus.

Paper	Year	#Tokens	#Tags	Link of the corpus
Bridges et al. [196]	2013	853,560	73,964	https://github.com/stucco/auto-labeled-corpus
Kim et al. [55]	2020	498,000	15,720	http://github.com/nlpai-lab/CTI-reports-dataset
Wang et al. [201]	2020	6574 (sentences)	36,412	https://github.com/SCreaMxp/DNRTI-A-Large-scale-Dataset-for-Named-Entity-Recognition-in-Threat-Intelligence
Alam et al. [75]	2022	106,991	4350	https://github.com/aiforsec/CyNER
Wang et al. [181]	2022	260,134	39,565	https://github.com/wangxuren/APTNER

encountering a term like Microsoft Word, some annotators treated Microsoft and Word as separate entities, classifying Microsoft as an organization and Word as a software name. In contrast, other annotators regarded Microsoft Word as a single entity and categorized it as a software name. To address these disagreements, the authors conducted a subsequent round of annotation. They provided additional training to annotators and redefined entity categories, merging certain entities into a single type.

In the majority of NER studies, researchers create NER corpora, but few of them choose to share these valuable resources with the public. Wang et al. [201] have contributed a cybersecurity NER dataset in which they gathered threat reports from various sources such as GitHub, government agencies, and security companies' Web sites. In the DNRTI dataset, they meticulously annotated 13 categories, including features, vulnerabilities, methods, organizations, industries, geographical areas, motives, timestamps, tools, security teams, sample files, attack types, and hacker organizations. Their annotation process employed the BRAT and utilized the widely accepted BIO labeling mode. However, the description of how they validated the quality of the annotations is not provided in sufficient detail. Table 11 lists the publicly accessible NER corpus.

8.2. Event identification

To respond rapidly to potential cyber threats, security analysts and IT personnel must stay informed about critical security events, regardless of their frequency of reporting. Given the rising number of cybersecurity incidents reported in articles, the capability to identify these events becomes a critical necessity. In NLP, event detection revolves around the discovery of details concerning entities within events, their respective functions, and the temporal and spatial attributes associated with these events. Event extraction is the process of extracting semantic and structural information from text, represented through typed phrases comprising trigger words and associated arguments [202]. It generally encompasses two primary sub-tasks: type classification, responsible for identifying specific event types within sentences, and element extraction, which captures trigger words and related arguments following diverse role patterns (schemas) to ensure comprehensive semantic understanding. When extracting events from the text, such as in the sentence “*Russian cybercriminals breached the International Criminal Court’s IT systems amid an ongoing probe into Russian war crimes committed in Ukraine*,” we can analyze both the event type and trigger words. In this context, the term “*breached*” serves as a trigger word, indicating an unauthorized intrusion or access to the IT systems of the International Criminal Court. This event is classified as a “*Attack.Databreach*” due to its involvement in compromising or accessing sensitive data, which is a defining characteristic of a data breach attack.

Trong et al. [203] generated a corpus for cybersecurity event detection, encompassing 8,014 event triggers spanning 30 event types and drawn from 300 articles. These event types were classified into four categories related to cyber vulnerabilities or attacks: IMPACT, ATTACK, PATCH, and DISCOVER. They also evaluated various models

for event detection using embedding techniques such as Word2Vec and BERT. The Document Embedding Enhanced Bidirectional RNN achieved 68.4% F1-score in the event detection task. In [81], the authors explored social media to extract cyber events and devised a method to prioritize potential cyber threats, considering criteria like user influence scores, named entity confidence, and keyword relevance. They mimicked the process of identifying named entities and keywords using TextRazor(NER API), and the TextRank [204] algorithm from Gensim. In addition, TextRazor furnished confidence scores for named entities, while TextRank assigned scores to keywords based on word graphs. In [180], Luo et al. addressed the extraction of cybersecurity event details from Chinese text, focusing on four event types: IncidentsOnVulnerability, RansomwareAttack, Phishing, and Data Breach. Their approach involved converting the event extraction task into a sequence labeling task. This process began by embedding characters and incorporating word information into character representations. They introduced a k-window-size BiLSTM to capture contextual information spanning sentences.

TCEDCL [205] is a cybersecurity event detection method that does not rely on event trigger word labeling but instead incorporates sentence semantics. It employs unsupervised comparison learning to find the most suitable candidate instances by using a contrastive learning model. They defined two main event types, further categorized into nine subtypes: VulnerabilityPatch, VulnerabilityDiscover, VulnerabilityImpact, SupplyChain, Malware, DDoSAttack, Ransom, Phishing, and Data Breach. A key benefit of TCEDCL is that it does not rely on detecting trigger words, simplifying the data sorting process during model training and lessening trigger word accuracy’s influence on the model’s performance. Instead, the system only needs to determine the event type, making it exceptionally flexible in adapting to new event types with minimal labeled data needed to train a new classifier. W2E(Words to Events) [206] identifies emerging cyber threats with low false positives and high coverage by monitoring individual words instead of using semantic clustering methods. They used NLP techniques, such as POS tagging, lemmatization, and NER, to reduce false positives effectively. Additionally, it filters and categorizes tweets based on keywords associated with five event types: Data Breach, DDoS attack, Vulnerability, Exploit, and Malware, while also monitoring CVE-related events separately. In [207] Mohammad et al. introduced a dual-level approach to cyber-event detection involving medium-level detection utilizing LDA and high-level detection based on Google Trends data. Their research encompassed comprehensive preprocessing steps, which included NER, POS tagging, symbol removal, stop-word elimination, and lemmatization. Also, they leveraged Word2Vec for feature vectorization and t-distributed stochastic neighbor embedding (t-SNE) for dimension reduction. This method achieved 95.96% accuracy in cyber event detection.

While existing methods can identify emerging cyber events, they often fall short of providing detailed threat characteristics. To overcome these limitations, Associated EE (AEE) [202] is introduced, which is capable of performing event extraction for every sentence in a document. AEE consists of two components: event type classification and element extraction, both utilizing BERT to grasp sentence-level semantic understanding. To classify event types, it constructs a document-level graph

that links sentences of diverse types and words within each document, utilizing a Document-Aware Graph Attention Network (DGAT) to represent sentences and grasp contextual information. AEE also introduces a novel schema to facilitate associations among argument roles and incorporates type-aware parameter inheritance to enhance extraction precision and enrich contextual knowledge across different event types and roles. Satyapanich et al. [208] introduced a framework aimed at classifying cybersecurity incidents relying on semantic schemas, focusing on five distinct event categories: Attack.Databreach, Attack.Phishing, Attack.Ransom, Discover.Vulnerability, and Patch.Vulnerability. Within this schema, there are 20 argument types designed to represent various elements and characteristics within these events. The study utilized two distinct word embedding methods, namely context-free (Word2Vec) and context-dependent (BERT). Word2Vec was trained on a set of 5,000 cybersecurity news articles, generating domain-specific embeddings with a size of 100. Conversely, BERT word vectors were derived by averaging Word-Piece embeddings from the fourth-to-last layer. The paper introduces Context2Vector [209], which employs representation learning for detecting event contexts to alleviate triage pressure. This enables security analysts to concentrate on resources requiring urgent attention. In essence, context represents a sequence of events generated from various perspectives, indicating specific security events. Context2event comprises four primary components: context extraction, context embedding, deviation detection, and human-in-the-loop annotation. To achieve this, Context2Vec utilizes LDA2Vec for creating event, context, and topic representations within the same vector space. Additionally, it employs the Local Outlier Factor algorithm to identify behavioral deviations and anomalies in dynamic contexts. Ultimately, this proposed approach generates informative and interpretable labels for human inspection, enhancing high-risk event detection models in dynamic environments. In [210], a method was introduced to predict future cybersecurity events employing probabilistic soft logic rules. This approach involved detecting explicit event references within a cyber attack kill chain by combining lexical matching and NER. Finally, in [183] the authors treat each vulnerability as an event and propose an approach called VE-Extractor, which is based on BERT Q&A, to automatically perform vulnerability event extraction from textual descriptions in vulnerability reports, including extraction of event trigger (cause) and arguments.

8.3. Relation extraction

Cyber Threat Intelligence (CTI) systems must aim to provide security experts with actionable information for detecting and mitigating cyber attacks. Traditionally, security experts have summarized this information in the form of IoCs, also known as low-level indicators, which include IP addresses, domain names, and malicious file hashes used to create detection rules. However, low-level IoCs have a short lifespan and lack the semantic context necessary for real-time attack detection. Additionally, attackers can easily modify such IoCs through obfuscation techniques, alterations to malicious code, changes in attack IP addresses, and the updation of phishing domains.

Collecting a substantial volume of valuable information from sources like social media, hacker forums, APT reports, threat reports shared by security companies, technical blogs, and more can transform this data into meaningful CTI. However, enterprises and organizations encounter three significant challenges when dealing with the vast, unregulated landscape of cyber threat information online. Firstly, CTI resources are complex, unstructured, and continually evolving. This complexity makes it challenging to extract actionable insights from the data. Secondly, there is a need for automated systems capable of synthesizing information about threat entities, attack semantics, and malicious actions. Automating this process can help rapidly identify and understand emerging threats. Lastly, organizations face the challenge of integrating threat data from multiple sources. They must harmonize data from various origins and understand the interactions

between different threat elements, all while discerning and filtering out ambiguous or duplicate information.

To effectively detect cyber attacks, trace the source of attackers, and understand their behavior, the fundamental requirement is extracting cyber objects such as attackers, tools, infrastructure, files, and more. This can be accomplished by extracting named entities and subsequently inferring relationships between these entities. Specifically, a significant body of research work can be categorized into two approaches: (i) *Pipeline Approach*: In this method, entities are extracted first, and ML or DL algorithms are employed in a subsequent phase to discern relationships between these entities. The information extracted includes entity heads, relations, and entity tails, providing crucial insights into various aspects such as the details of the attack incident, attacker behavior, threat actions, exploitation of vulnerabilities, and threat hunting. However, a critical drawback of this approach is that errors introduced during the entity extraction phase can propagate and influence the accuracy of relationship extraction, (ii) *Joint-Entity Relation Model*: In contrast to the pipeline approach, the joint-entity relation model involves learning entities and their relationships simultaneously. This approach aims to capture associations between entities and their corresponding relationships in a more integrated manner, potentially mitigating the impact of errors in entity extraction on the subsequent relationship extraction stage.

By analyzing the related literature about Relation Extraction (RE) from CTI, we can identify several common components that such solutions consider and design. In general, we can summarize the task of extracting relationship triplets as shown in Fig. 8. In the following paragraphs, we will analyze in detail these components by highlighting the different approaches adopted.

Tactics Techniques & Procedures (TTP): SeqMask [211] uses *Multi-Instance Learning (MIL)* to find TTPs in CTI. It predicts TTP labels using behavioral keywords and conditional probabilities. SeqMask has two mechanisms to ensure credibility: expert validation and blocking existing keywords to assess their effect on accuracy. Experiments show SeqMask achieves a notable 86.07% F1 score for TTP classification and improves TTP extraction from extensive CTI and malware data. The *Attention-based Transformer Hierarchical Recurrent Neural Network (ATHRNN)* [212] performs TTP extraction from unstructured CTI. This involves several key steps. First, they design a Transformer Embedding Architecture (TEA) to capture high-level semantic representations of both CTI and the ATT&CK taxonomy. Then, they develop an Attention Recurrent Structure (ARS) to model the relationships between tactical and technical labels within ATT&CK. Lastly, they create a joint Hierarchical Classification (HC) approach to understanding the interdependencies between tactics and techniques, enabling the prediction of the final TTPs. ATHRNN demonstrates superior performance in extracting TTPs, surpassing benchmarks like TF-IDF, Word2Vec, SVM, Deep Neural Networks (DNNs), BERT, and Text-CNN. Recently, the authors proposed TTPpredictor [213], which maps CVE descriptions to threat actions mentioned in the MITRE knowledge base. This work addresses two critical issues: (1) the scarcity of datasets containing labels for threat descriptions and (2) the semantic disparity between threat descriptions and MITRE techniques. To tackle these issues, the authors created a large corpus of threat reports mentioning CVEs correlating with MITRE descriptions. They estimated the correlation by matching the verb-object pairs of the corpus with the MITRE text using semantic role labeling. This resulted in a labeled dataset, which the authors used to fine-tune BERT to predict actions. Additionally, they compared the performance of TTPpredictor to ChatGPT by injecting 20 CVEs and prompting both tools to predict MITRE techniques for a vulnerability description. They found that ChatGPT failed to predict the techniques for vulnerability reports 80% of the time.

HMCAT [214] framework, integrates the augmentation of data, and operates in two key stages. In the Pre-processing Step, data augmentation is performed using LLMs to generate new CTI samples based on instructional prompts and exemplars. To ensure data quality, a

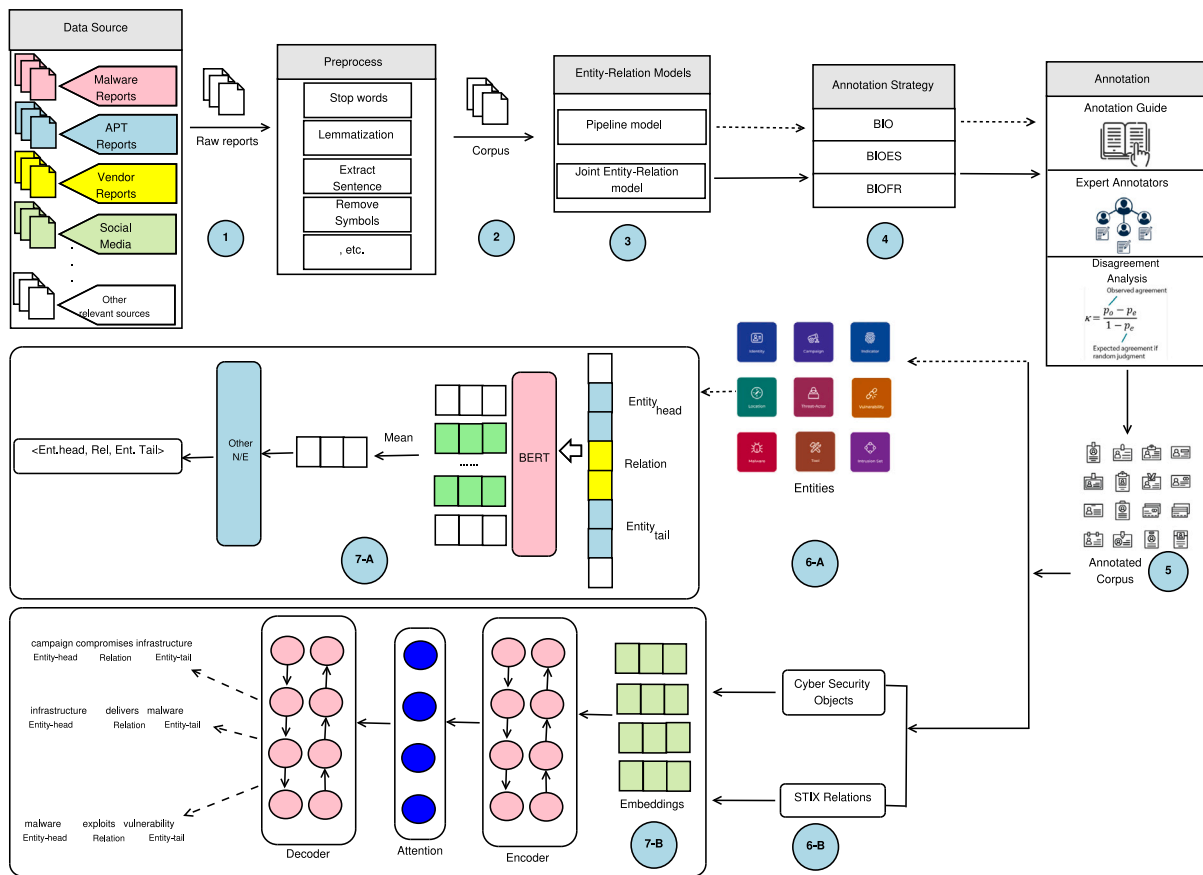


Fig. 8. Relationship Extraction Pipeline.

critic model (SBERT) filters out low-quality samples, such as those with mismatched attack techniques, redundant information, or irrelevant text. Additionally, IoC entity extraction is performed using regular expressions and NER models. In the Training Step, extracted text descriptions, IoC entities, and attack techniques are encoded in vector representations. Then a hierarchy-aware mapping method is used to map relationships within the label hierarchy, improving classification accuracy. This structured approach enables more effective mapping of unstructured CTI to ATT&CK techniques, improving the automation of threat intelligence analysis and attack attribution.

Threat entity extraction and relation extraction: CDTier [72] identified five types of threat entities and 11 types of relationships between these entities based on standards like STIX and real-world CTI usage. The process began with manually annotating a text corpus to create a labeled dataset. This dataset was then used to train a word2Vec model, which represents word meanings using vectors for word semantics, position, and grammatical relationships. These vectors form the basis for a neural network. After further feature extraction, the significance of semantic features was enhanced using Softmax, resulting in the extraction of entity-relationship pairs. To extract semantic information from cybersecurity documents, the authors [215] used Java Annotation Patterns Engine (JAPE) with General Architecture for Text Engineering (GATE) framework. The approach focused on noun and verb tokens that describe the relationships between entities, but the accuracy was only 32.69% due to the limited number of samples.

The CARE (Cyber Attack Relation Extraction) [216] system extracts entities and relations from Chinese cybersecurity Web sites. CARE comprises four important phases: (i) A crawler collects text from several Chinese cybersecurity Web sites, (ii) Google Translate converts each article to English, (iii) Entity and relation extraction of nine STIX and ATT&CK and five relationships, (iv) Finally, a pre-trained BERT model fine-tunes with labeled data to capture relationships between entities

in the articles. The evaluation of 150 articles using BERT, RoBERTa, and ALBERT demonstrates that BERT outperforms its counterparts, showing superior performance. The paper [217] proposed an information extraction process that involves two aspects. First, it employs the Dictionary Template-BERT-BiLSTM-CRF (DT-BERT-BiLSTM-CRF) entity extraction model with a CRF to obtain entity labels. It creates a comprehensive dictionary template containing 3857 entries of entities, such as hacker groups, malware names, attack types, and security teams. The dictionary template is used to enhance the accuracy of the labels predicted by CRF. Second, the relationship extraction model is based on distant supervision and reinforcement learning, with one of its essential components being the Piecewise CNN classifier. To improve the quality of instance/sentence selection, reinforcement learning assigns rewards to the sentences based on the F1 score. This rewards function also improves sentence selection. Research on relational triplet extraction for IoV threat intelligence is gaining attention, though existing literature remains limited. The proposed approach employs a Span-based multivariate information-aware embedding network [218] to comprehensively analyze multiple information aspects within a sentence. By leveraging Deep Neural Networks and attention mechanisms, it fuses multivariate information into sentence embeddings, enhancing contextual understanding. Additionally, a Type-Aware Graph Convolutional Network (TA-GCN) is introduced to strengthen token relationships, effectively capturing critical information for improved threat intelligence representation.

The author in [176] introduces CyberRel, a *joint entity and relation extraction model* for cybersecurity concepts. CyberRel tackles the problem by performing sequence labeling, generating unique labels for relations that capture crucial information about the entities and their roles. CyberRel uses a pre-trained BERT model to create word embeddings, subsequently leveraging a BiGRU neural network with an attention mechanism to extract semantic features. Lastly, combining

BiGRU and CRF constructs cybersecurity triples. CyberRel achieves an F1 score of 80.98%, surpassing the performance of the conventional pipeline model. Recently, the authors proposed the CTI-JE [219] model which treats joint extraction as a table-filling problem to recognize overlapping entities. They extract CTI from Microsoft's threat intelligence website, defining 9 types of entities and 5 types of relations based on the STIX standard. To address BIO tagging limitations, they used span annotation to record entity positions. SecBERT serves as the encoder, with separate MLPs for NER and RE, along with a shared representation to enhance interaction. They incorporate relative positional encoding, dilated convolution, and channel attention to capture contextual dependencies. Results show that CTI-JE outperforms baseline models in extracting CTI triplets. Further study in [220] extracts cyber relation triples using a joint extraction technique, addressing pipeline method issues. The BIEOS tagging scheme enables joint tagging and RE on CTI data. The RoBERTa-BiGRU-CRF model processes text, where RoBERTa generates embeddings, BiGRU predicts word probabilities, and attention and CRF layers enhance tagging accuracy. Relations are extracted using a relation-matching technique, leveraging knowledge graphs. Entity embeddings from RoBERTa and GCN are combined in a feedforward neural network to compute a relation-matching score, which is merged with a confidence score to select the final predicted relations.

The work in [221] focuses on extracting relational triples from CTI data using an Enhanced Representation and Binary Tagging Framework (ERBTF). The approach begins by integrating relation embeddings with word embeddings to construct a richer initial hidden representation. To refine contextual feature extraction, a dilated convolutional encoder is employed, which incorporates a dilated Convolutional Neural Network (CNN), a gating mechanism, and residual connections, to enhance the representation of contextual dependencies. Furthermore, an attention module, combining multi-head self-attention and a position-wise feed-forward network, is utilized to emphasize words that are crucial for determining specific relations. Finally, a binary entity tagging mechanism is applied to identify subject and object entities across different relations. The authors [222] highlight two fundamental challenges in joint entity-relationship extraction methods. First, existing approaches treat both subtasks – entity recognition and relationship extraction – using identical embeddings. These methods typically rely on pre-trained models as encoders, applying the same embeddings for both tasks, which overlooks fine-grained feature representations specific to each. Second, conventional techniques classify relationships based on vector representations of entity pairs, failing to capture meaningful feature interactions during the relationship extraction phase. To overcome these limitations, authors propose CTI-TFN, a task-specific Fourier Network-based (FNet) model for joint entity-relationship extraction. Their approach is evaluated on a dataset of 4,000 sentences scraped from Microsoft's Threat Intelligence website. The framework employs two distinct Fourier networks – one dedicated to entity recognition and the other to relationship extraction – enhancing task-specific feature learning. Additionally, a gated attention mechanism facilitates interaction between these subtasks, ensuring a more cohesive extraction process. This study focuses on extracting nine entity types and five relationships aligned with the STIX 2.0 standard, contributing to improved structured threat intelligence extraction.

Vulnerability detection, and exploitation: Predicting the exploit potential of vulnerabilities is vital for decision-makers to prioritize their actions and address critical flaws. However, the limited size of the vulnerability descriptions makes it challenging to train a comprehensive NLP model. ExBERT [76] is an enhanced BERT model tailored for *exploitability prediction*. The approach involves word piece tokenization for embedding extraction, addressing infrequent cybersecurity terms and polysemy. Furthermore, pre-trained BERT is fine-tuned using a domain-specific dataset. Additional layers, including a Pooling Layer and Classification Layer, are added to the fine-tuned BERT model to capture sentence-level semantic features and predict vulnerability

exploitability using an LSTM model. ExBERT achieves an impressive 91% accuracy in exploitability prediction.

The author proposed VE-Extractor [183], a tool for *vulnerability analysis* and *classification* that uses dynamic information from text. VE-Extractor was trained on 4638 vulnerability reports, and it can extract vulnerability events, which consist of vulnerability triggers (causes) and arguments (such as attacker, consequence, operation, location, and version). The triggers and arguments were first manually extracted using a tagging scheme called BIOFR. Then, the tool used BERT-BiLSTM-CRF to map the descriptions to triggers. Next, FastText was used to classify the triggers into different types. Finally, BERT Q&A was used to design question templates to extract the vulnerability arguments, depending on the trigger.

The paper [223] proposes an unsupervised approach to label and extract important *vulnerability concepts* (e.g., root cause, attacker type, impact, etc.) from vulnerability descriptions. The system first constructs a vocabulary of absolute paths and relative paths, which captures the syntactic differences between sentence structures and phrase expressions. Then, a Categorical Variational Autoencoder (CaVAE) is used to learn embeddings for each path. CaVAE places identical paths in the same cluster, representing the same type of concept. Finally, vulnerability concepts are extracted based on the human labels and considering noun phrases.

The paper [82] proposes a Vulnerability Information Extraction Model (VIEM) to detect inconsistencies in vulnerabilities reported in NVD and CVE, particularly related to software names and versions. Security analysts primarily rely on vulnerability reports to protect their organizations from ongoing and future cyber attacks. However, inconsistencies in the information published by different sources can be a concern, as specific versions of unpatched software expose stakeholders to more significant security risks. VIEM utilizes the NER and Relationship Extraction (RE) models to address this limitation by recognizing software names, versions, and associated relationships between them. VIEM's NER model encodes the text sequence into word and character-level vectors. It leverages bi-directional GRU to predict one of the following labels for each word: (i) software name, (ii) software version, or (iii) others. Subsequently, the RE model takes the entity labels predicted from the NER model and creates all possible combinations of software names and versions. It encodes this combination as one-hot vectors. Finally, the RE model takes a sequence of vectors as input and, using a hierarchical attention mechanism predicts if the pairing of the software name and version is appropriate.

Cyberhunters and security analysts rely on diverse sources to comprehend threats and formulate mitigation strategies. However, vulnerabilities originate from one document, while exploitation and countermeasures exist in others. Thus, manually establishing connections between documents poses a challenging and intricate task. The work [224] identifies available entity information in one document and deduces missing data from other documents using NLP techniques. It accomplishes this by analyzing free-form text to infer concealed relationships that illustrate threat actions and mitigation strategies. The investigation was conducted on the BRON dataset. Text is subsequently represented in free form using four Language Embedding Models (LEM): Bag-of-Words, GloVe, BERT, and Fine-tuned BERT. Following this, a Random Forest classifier was trained on a dataset containing labels denoted as links and no-links. GloVe and BERT excelled in illustrating relationships, while Bag-of-Words yielded unsatisfactory results. However, GloVe yielded mixed conclusions. The proposed approach [225] employs a probabilistic framework, specifically statistical relational learning, to detect product weaponization. Understanding the connections between products, vulnerabilities, and exploits is beneficial to various stakeholders, including vendors, users, and product operators. Vendors can leverage this information to alert stakeholders about critical risks and emphasize the importance of patches. Users can utilize these insights to determine whether to prioritize or delay patching. Meanwhile, operators can analyze missing relationships to enhance

attack mitigation strategies. The method aggregates data from diverse cybersecurity sources (e.g., Shodan, NVD, and ExploitDB) to establish the relationship between products and exploits. This is achieved by leveraging the known associations between vulnerabilities and exploits (sourced from CVE and ExploitDB) and between products and vulnerabilities (gathered from CPEs and CVEs). Additionally, findings demonstrate that relational regression trees can extract meaningful insights from the model, incorporating factors such as product types, vendors, and versions. TRACE [226] extracts causal factors associated with adversarial techniques and generates causal graphs to reveal interdependencies within CTI reports. In particular, it maps attack techniques to their underlying Common Weakness Enumerations (CWE) and vulnerable data components to uncover intricate relationships between adversarial tactics and system vulnerabilities. Leveraging SBERT and BiLSTM, TRACE automates the extraction of text related to causal factors from CTI reports, enhancing the efficiency of cyber threat analysis.

Action/Event extraction and detection: CASIE [90], a cybersecurity event extraction system, detected events in 1000 English news and CTI texts by using a combination of the Attention Mechanism, BiLSTM model, and linguistic features. CASIE defined five event types, their corresponding semantic roles, and the arguments that fulfill these roles. However, CASIE defined five event types but only focused on event arguments near the event triggers. This approach overlooked event arguments located further away from the event triggers despite events often spanning multiple sentences within a document. Husari et al. [227] developed ActionMiner, which utilizes entropy and mutual information for threat extraction. Specifically, these information-theoretic measures are employed to identify object-verb pairs representing threat actions, such as “execute-file”, “run-registry”, “delete-file”, and so on. ActionMiner reported improved performance compared to alternate solutions employing the Stanford dependency parser solution. Also, ActionMiner faced limitations in identifying actions from CTI text containing syntactic blocks composed of verbs, verb-direct objects, verb-indirect objects, verb-subject complements, etc.

CyEvent2vec [228] is an event-embedding framework based on Heterogeneous Information Network (HIN). The framework aims to predict events using three datasets comprising three actions: malware, misuse, and physical. The framework creates an event matrix to depict the event with attributes, where the matrix elements denote the relationship between the event and the object with attributes. Subsequently, using an autoencoder, intricate relationships between the heterogeneous objects are obtained, which produces low-dimensional embeddings. This autoencoder comprises many sub-encoders, each designed to generate embeddings for the object or attribute. Finally, to map the event matrices to the embedding space, hidden layer embedding of sub-autoencoders is combined.

Ex-Action [229] extracts threat actions from CTI reports by applying rule-based matching and multimodal learning. It identifies each threat action as a triple (subject, object, verb) obtained by matching syntactic rules. The subject and object correspond to security entities, and the verb describes the semantic relationship between the entity pair. The TF-IDF and BM25 algorithms calculate the similarity between candidate actions and a CTI report. Ex-Action also identifies threats using multimodal learning algorithms. To evaluate the completeness of the threat action obtained using the proposed method and manual extraction, Ex-Action utilizes Normalized Mutual Information (NMI). Experiments on 243 CTI documents from ATT&CK demonstrate the efficacy of the proposed solution, which provides 79.09% accuracy and 85.26% NMI.

Malware behavior : Huang et al. [230] used the MITRE website to develop MAMBA, a Malicious Behavior Analysis system. They employed DL and ATT&CK knowledge to uncover TTPs and create behavior profiles for Windows malware. MAMBA has three phases (a) Extraction: Knowledge from ATT&CK is combined with malware traces from sandbox execution, (b) Fusion: ATT&CK info and execution traces are merged using PV-DM embedding and (c) Threat ID: Malicious

behaviors (TTPs) in malware traces, involving API calls, are identified. The model achieved over 90% precision and recall.

CTDroid [231] employs NLP technologies to comprehend the behavior of Android malware as documented in security blogs, thereby enhancing the detection and classification of such malware. The proposed approach specifically addresses the semantic gap between natural language and programming language. Initially, textual data is gathered from technical blogs published between 2011 and 2017, and sensitive behaviors, represented as subject-object pairs, are extracted through Stanford dependency parsing. Subsequently, the frequency of these sensitive behaviors is estimated. Next, each malware sample undergoes reverse engineering using an open-source tool, and permissions, API calls, and intents are extracted from the AndroidManifest.XML file. Descriptions corresponding to these attributes are collected from Android’s official documentation. The similarity between the frequently occurring behaviors and the behaviors extracted from the malware samples is then calculated. To assess the similarity of the textual descriptions with the Android system, words are transformed into word2Vec representations. Furthermore, the matching of sensitive behaviors extends to the code level, where keywords within function definitions are extracted, and similarity is computed as previously described. Finally, each sample is represented as a vector, utilizing sensitive behavior frequencies and a boolean vector to train five ML models. In conclusion, CTDroid achieves an impressive true positive rate of 95.8% with only a 1% false positive rate.

Analyzing IoC Relationships: In their study [77], the authors developed a trigger-enhanced system (TriCTI) to determine the *relationship between IoCs and campaign stages* from unstructured cybersecurity reports spanning 21 years (2000–2021). They successfully extracted campaign trigger phrases indicating six different campaign stages. The authors employed the BERT model to train embeddings for sentences, trigger vectors, and IoCs simultaneously. This approach was used to classify sentences into their respective campaign stages. The scarcity of data during the training stage was addressed by augmenting sentences using CBERT. The classification model reported an accuracy of 86.99%. However, the model misclassified some examples that contained multiple campaign stages. In [232], authors proposed a system to analyze risk in an Additive Manufacturing (AM) system. The framework recognizer module identifies domain-specific IoCs, which include values such as *organization_domain_specific*, *regional_source_specific*, and *regional_target_specific* from multiple sources. Subsequently, IoCs are modeled using a Heterogeneous Information Network (HIN), which consists of meta-paths and meta-graphs. HIN captures eleven semantic relations amongst IoCs. The proposed method uses similarity measures on meta-paths and meta-graphs to determine access to the threat. The system also provides a ranking of IoCs based on diverse factors such as lifetime, number of vulnerabilities exploited, and the attack frequency.

APT Analysis: SECCMiner [233] analyzed the trend for the 13 most-mentioned techniques in its data from 445 APT reports collected from 2008 until 2017. The system used Context Free Grammar to extract all unique noun phrases and compute the TF-IDF score for each noun phrase. These TF-IDF vectors were used to determine the most identical document that matches the 13 most common APT techniques listed in the MITRE ATT&CK knowledge base. Furthermore, SECCMiner conducted an APT trend analysis to report that PowerShell scripts are one of the rising techniques, using waterhole together with spearphishing and correlating zero-day exploits with Adobe Flash software. Recently, authors proposed ContextMiner [234], a framework that extracts contextual features using a combination of part-of-speech tagging, dependency parsing, and syntactic grammar. It was evaluated for two aspects: retrieval of contextual phrases from labeled and unlabeled APT corpora, and its effectiveness as a pre-processing step for NER. ContextMiner reduces the size of the original dataset by 70%, reducing computational and storage requirements while producing the highest performance in NER.




Fig. 9. An example cybersecurity knowledge graph.

8.4. Knowledge representation using graphs

Cybersecurity knowledge graphs are collections of nodes and edges that carry semantic information about vulnerabilities, threat actions, attack patterns, etc. Entities such as attackers, product names, organizations, tools, and malware are represented as nodes, and the relationships between entities are represented through edges (see Fig. 9). Modeling textual information from different security articles as a KG offers three advantages: (i) security information from heterogeneous sources can be integrated into the KG to provide causal inference, entity correlation analysis, or semantic reasoning between the concepts, (ii) knowledge can be depicted in a structured form and visualized in a graphical format, and (iii) security analysts can query, reason on existing facts to validate consistency in data, derive novel patterns, or refine the KG by inputting new relations to improve semantic modeling.

To construct a robust cybersecurity KG (see Fig. 10), we can begin by leveraging pre-designed cybersecurity ontologies developed by esteemed security organizations or cybersecurity experts. These ontologies are tailored for specific applications such as intrusion detection, vulnerability exploitation, and malware categorization. They comprehensively define the concepts and intricate relationships between various attributes. The subsequent step involves Information Extraction (IE), wherein the objective is to meticulously extract cybersecurity-named entities and their corresponding relationships. This extracted data then forms the foundation for constructing the KG, wherein entities are represented as nodes, and relationships are depicted as edges. Following the creation of the initial KG, it is imperative to execute entity linking and disambiguation processes to meticulously eliminate the propagation of errors and enhance the graph's precision. Moreover, to continually enhance the quality of the KG, we should actively seek timely feedback from security experts. This iterative feedback loop enables us to refine and enrich the graph with the latest insights and developments in the field of cybersecurity.

The authors in the paper [235] generated a repository of CTI concerning various medical devices and their known vulnerabilities. They augmented this repository by incorporating data from sources such as Wikidata and information regarding insecure medical devices from the Global Unique Device Identification Database (GUDID). Using this augmented repository, they created a Cybersecurity Knowledge Graph (CKG). The proposed methodology enhanced the knowledge of entities and relationships derived from the Unified Cybersecurity Ontology (UCO) to represent CTI within the CKG. To devise a means of expressing node features that proved valuable in predicting their behavior within specific neighborhoods of given nodes, the authors employed the Breadth-First Search (BFS) algorithm. They chose BFS because it provides a local, microscopic graph perspective. Once they defined the n -neighborhood for each node, they used the node2vec algorithm to generate embeddings. As a result of their efforts, they reported a notable 31% increase in the Mean Average Precision (MAP) value when evaluating an information retrieval task. To enable collaboration in detecting and preventing attacks, organizations must store and represent information in a structured format. KAVAS (Knowledge-Assisted Visual

analytics) [236] represents STIX Domain Objects and their relationships in a graph database, specifically Neo4j, and connects this database to the visual interface. This interface aids security analysts in *comprehending and analyzing threat incidents*. KAVAS also facilitates externalization, allowing security analysts to enhance their knowledge by modifying existing threat descriptions.

The AI system Cyber-All-Intel [237] aids security analysts in knowledge extraction, representation, and analytics. It utilizes a 'VKG structure' that merges word embedding and KG into a vector representation. To extract information from the KG, it employs SPARQL. LADDER [238] is an automated framework that extracts TTPs from CTI sources related to malware and APTs. This information is organized into a KG using an ontology and TTPClassifier. The framework enables predictive analysis, threat hunting, and attribution of APT groups. TTPClassifier employs a novel ML algorithm to extract TTPs from CTI reports, including IoCs and TTPs, and categorizes them using standardized MITRE ATT&CK pattern IDs. Satvat et al. [83] developed a methodology to extract comprehensive insights into *attack behavior* from CTI reports. Then, this information was organized for identifying APTs and conducting threat hunting. They actively transformed intricate descriptions in CTI reports into a provenance graph, with nodes representing entities and edges denoting system calls. They actively employed various NLP techniques such as co-reference resolution, conversion from passive to active voice, text summarization, and word standardization to achieve this transformation. Furthermore, they actively used Semantic Role Labeling to extract attack behavior elements, including subjects, objects, and actions, ultimately facilitating the conversion of unstructured text into structured provenance graphs. Pingle et al. [239] proposed a solution called RelExt, a Feed Forward Neural Network (FFNN) classifier, to extract semantic triplets from cybersecurity text and develop KG to *detect cyber threats*. RelExt first removes semantic triples that do not follow the semantics of STIX2.0 and entities that are not present nearby. It then predicts the relationship between the remaining entities using the FFNN classifier. RelExt achieved an accuracy of 96.61% in identifying relationships between entities and successfully predicted 700 relationships from Dark Caracal and CrossRat malware threat reports. Modeling CTI is challenging due to the heterogeneity of cyber-threat infrastructure nodes (i.e., domain name, IP address, malware hash, and email address). To address this challenge, the authors proposed HinCTI [240], a system that models CTI on a heterogeneous information network (HIN). HinCTI integrates infrastructure as nodes and the relationships between them as edges. HinCTI uses meta-schema, meta-paths, and meta-graphs to capture the semantic interdependence between infrastructure nodes. *To identify the threat types* of infrastructure nodes involved in CTI, the authors define a measure of similarity between threat infrastructure nodes based on meta-paths and meta-graph instances. HinCTI also presents a heterogeneous Graph Convolutional Network (GCN) approach that uses this measure to identify the threat types of infrastructure nodes.

The authors introduced CSKG4APT [178] as an APT KG for attributing cyber threat attacks. CSKG4APT's knowledge graph finds application in *threat hunting* and *intrusion analysis* following the diamond model. It defines 12 knowledge types, such as attackers, attack types, tools, assets, and more. The knowledge types/entities from the bilingual threat intelligence corpus are extracted using BERT-BiLSTM-GRU-CRF. Semantic relations between these entities are established by drawing from the STIX framework and NVD.

AttacKG [241], which automatically extracts attack behavior from multiple reports to reduce human effort and speed up response time. Specifically, the authors *summarize techniques* into Technique Knowledge Graphs (TKGs). They first parse CTI reports from different sources to extract attack-relevant entities (using NER) and their interdependencies. Next, they construct technique templates from the technique procedure examples extracted from the MITRE ATT&CK knowledge base. These technique templates are also represented as graphs. Finally, they propose aligning the multiple subgraphs built from the CTI sources

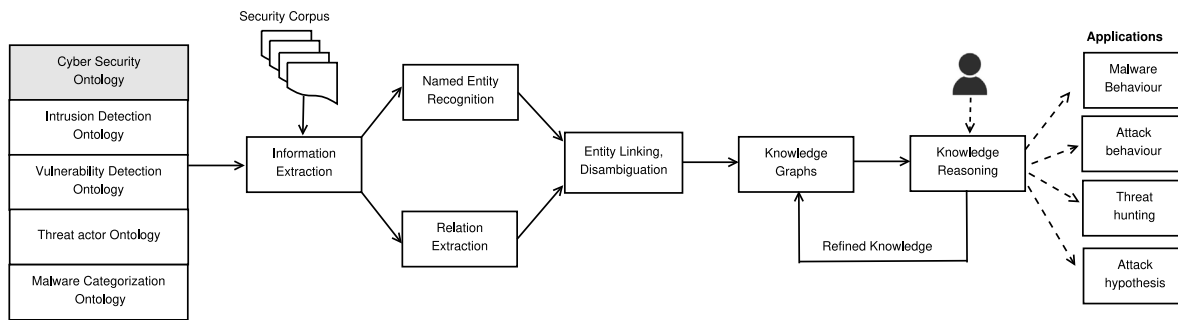


Fig. 10. Building Knowledge Graphs.

with the technique template graphs. They ascertain the alignment of graphs by computing (i) node alignment score and (ii) graph alignment, which determines the overall alignment score of CTI graphs with a technique template graph. The alignment of the attack graph and corresponding technique template graph enhances the attack graph with knowledge from templates and simultaneously updates the template with rich information from multiple CTIs. AttackKG accurately constructs attack graphs from reports, with F1 scores of 0.887 and 0.896 for entities and dependencies extraction. Recently, the authors presented the Open-CyKG framework in [156], comprising three modules: the extraction of relation triples from APT reports, NER module for labeling cybersecurity entities, and the construction of a KG. The process begins with the initial parsing of textual documents. Embeddings corresponding to words, Parts of Speech (POS), and predicates are inputted into a Bi-GRU for performing sequence labeling. The output is then passed to an additive attention layer, which assigns a higher weight to essential words in the text. Subsequently, a softmax layer outputs individual probability distributions covering different tags. The NER module further employs Bi-GRU followed by Dense layers and Conditional Random Fields (CRF) to label each entity in the cybersecurity corpus. Finally, the KG is created using the relation triples, with each node augmented with the named entity tags. The KG is disambiguated using hierarchical agglomerative clustering, and ultimately, the graph is visualized using Neo4j. The Hyper Attack Graph [242] framework applies NLP methods, including a joint extraction model with a multi-head mechanism, to accurately extract security entities and their relationships from CTI reports. These extracted elements are then integrated into a hypergraph structure to represent the complex relationships that allow visual representation of the tactics and techniques involved in cyber attacks. Recently Wang et al. [243] developed KnowCTI, integrating graph attention networks with entity and relation extraction to construct a threat intelligence graph. They treat entity extraction as a sequence labeling task and relation extraction as a classification task, leveraging a knowledge ontology to extract entities and relations. A cybersecurity-aware embedding method refines knowledge triples using an attention and gate mechanism, forming a sentence tree processed as a graph for contextual integration. Using BERT-based embeddings, the model achieved 90.16% F1 for entity extraction and 81.83% F1 for RE on OpenCVE, Malvuln, and ATT&CK datasets, demonstrating its effectiveness in structuring threat intelligence from unstructured data.

The paper presents an automated method, K-CTIAA [244], for analyzing CTI. K-CTIAA retrieves cybersecurity terminology information from the KG and inputs it into the pre-trained model. Additionally, K-CTIAA offers countermeasures to aid security experts in defending against incoming cyber attacks. K-CTIAA consists of three modules: pre-processing, CTI analysis, and countermeasure modules. It converts CTI articles from various sources into a KG. It extracts named entities from query sentences to generate triples, matching pairs with the KG is also called knowledge query. The knowledge query is subsequently transformed into a sentence tree, augmenting each token with additional information extracted from semantic triples in the KG. The sentence tree is further converted into a linear structure using Depth-First (hard

coding) and Branch-First (soft coding) approaches. Encoding is accomplished using BERT, considering position index via soft coding and tokens extracted through hard coding. The visible matrix in this work serves as self-attention, offering precise contextual information about the sentence. Finally, the countermeasure module suggests countermeasures based on the KG created from the ATT&CK knowledge base. In [245], authors proposed a framework called TIMFuser that scraped content from security pages, employing a breadth-first search technique to ensure quality. The extracted texts are preprocessed through normalization and removal of irrelevant content, followed by the use of the Longformer model for the contextual representation of long CTI reports. For recognizing important attack details, they used a fine-grained BERT-based model with a BiLSTM layer. Finally, TIMFuser constructs attack graphs using semantic role labeling to represent relationships between attack subjects and methods.

To help security analysts formulate *attack hypotheses* about the system's state, understand attackers' objectives, and prepare an action plan for responding quickly to an attack, the authors in [246] focused on detecting higher-level IoCs. Specifically, it aims to detect techniques and then deduce IoCs from observable artifacts. They introduce a tool called the Attack Hypothesis Generator (AHG), which extracts techniques from low-level telemetry data and refines these hypotheses using link prediction. The proposed approach initiates hypothesis generation, traversing a KG if the hypotheses involve techniques closely resembling the currently observable data (COD). It generates initial hypotheses using TF-IDF (utilized for scoring techniques), Naive Bayes (to assess the relevance of techniques to the COD), and Multi-layer Naive Bayes inference (to determine connected techniques). Subsequently, the hypotheses refinement module elevates the scores of techniques frequently employed together in launching an attack. Notably, this method relies on link prediction techniques and similarity measurement metrics. Consequently, when given a KG, the link prediction identifies missing edges (relationships) that are likely to occur.

Vulcan [247] serves as a CTI system designed specifically for ransomware attacks. It identifies cybersecurity entities through a custom NER process utilizing the Threat Entity Identifier (TEI). The TEI employs BERT, BiLSTM, and CRF models to transform the embeddings of input tokens into six labels that denote ransomware-related entities. Following the entity recognition step, the Entity Linker module performs two crucial tasks: (i) entity segmentation and (ii) entity integration. These tasks address potential issues where (i) the same entity might refer to different things or (ii) multiple variants of entities refer to the same thing. The relationships between entities are established using a pre-trained BERT model. Before training BERT, the input sentence is processed by adding special markers at the beginning and end positions of the entities. To determine the type of relation between two given entities, the RE module utilizes the final hidden state vectors of the BERT model for the [CLS] token and the two entity types. Vulcan stores the extracted CTI data in a graph database and offers security practitioners search APIs to discover the relationships of a given entity with other entities. In [96,195], the authors employed a pipeline approach to extract entities and relationships, subsequently constructing

a knowledge graph using Neo4j. In [195], SecureBERT with a BiLSTM layer and a Time-Distributed Dense (TDD) layer was used for entity extraction, while in [96], the authors utilized BERT with BiLSTM-CRF. For relationship extraction, the authors applied NLP techniques such as Attention-based Bidirectional Long Short-Term Memory (Att-BiLSTM) in [195] and BERT in [96]. Recently, the authors of [248] have developed a joint entity and relation extraction model for *threat intelligence extraction and fusion*. They fused entities extracted from various data sources by enhancing the Levenshtein distance algorithm, imposing an additional penalty on edit operations involving numbers rather than characters, to create a preliminary CKG.

Zhao et al. [177] proposed a HINTI framework to *model the intricate relationship between IoCs* applicable for real-world applications such as (i) profiling and ranking CTIs, (ii) model attack, and (iii) vulnerability analysis. The proposed method enables the creation of a CKG, which involves the following tasks: IoC extraction, relation extraction, and cyber event detection. They propose a multi-granular approach for the extraction of IoCs. The HINTI framework, using BiLSTM+CRF, transforms the words in textual documents into different n-gram types, including character n-grams, 1-grams, 2-grams, and 3-grams, and identifies IoCs of varying lengths and types. The authors also created heterogeneous graphs to obtain granular relationships between IoCs. These graphs were able to summarize nine relationships between six IoCs. Finally, the authors proposed a threat intelligence computing framework using a Graph Convolutional Network (GCN) for effective knowledge discovery. Shi et al. [249] developed a threat knowledge graph by consolidating information from common threat databases, including CVE, CWE, and CPE. They translated database entries and their associations into triples and embedded them into a vector space using models such as TransE, DistMult, and ComplEx, with TransE demonstrating the best performance. Their approach achieved notable results, such as a 0.606 Hits@10 score for CVE-CPE predictions. Furthermore, the predictive accuracy of the graph was improved by removing outdated entries and incorporating additional data from sources like CAPEC and CVSS vectors. In [250], the authors proposed a cyberattack detection method that utilizes knowledge graph reasoning. They gathered data such as vulnerability details, logs, and network traffic using a Turtle file parsing algorithm for data preprocessing. Also, they created a semantic framework by constructing an ontology, aligning system states with a security knowledge graph to support in-depth network data analysis. To handle asymmetric relationships in the data, they applied knowledge graph embedding (KGE) techniques such as the TransH model. Jaouhari et al. [251] leveraged NLP techniques to gather and process CTI data from diverse sources, focusing on identifying vulnerabilities and threats to IoT. They employed methods such as keyword matching, cosine similarity, and document summarization to analyze and filter relevant articles, which were then stored in a machine-readable format for further analysis and visualization within a knowledge graph using Gephi.

Extracting IoCs is a crucial step in thwarting cyber attacks, and it is equally vital for security operators to assess the “*maliciousness of IoCs*” to enhance decision-making and cut down on security operation expenses. To gauge the malevolence of IoCs, Kazato et al. [252] employed GCN, which exploits two aspects related to IoCs: (i) the distinct features of individual IoCs, encompassing statistical, network-based, and OSINT-based attributes, and (ii) the IoC graph. Specifically, GCN utilizes individual IoC attributes, adjacency data, and label information to determine the malignancy of an IoC. Their approach successfully identified 88.3% of malicious domains and achieved a performance of 96.3%.

SecKG [253] utilizes a KG to analyze event logs within a Windows system. The KG represents entities mentioned in the logs, such as processes, files, and registry entries, and the edges between the nodes depict actions between these entities. The detection module within SecKG employs rules expressed in first-order predicate logic to extract attack techniques from the KG. Subsequently, the system employs a

Graph Convolutional Network (GCN) to create graph representations of the training data, and the prediction module aims to extract subgraphs that match specific attack techniques.

Poirot [155] is a system designed to emulate kernel audit log records to encompass information flows within a system, serving the purpose of *threat hunting*. It comprises three crucial components: (a) the creation of a provenance graph, (b) the generation of a query graph, and (c) the alignment of these graphs. The kernel audit reports are represented as a provenance graph, with nodes representing entities such as files, processes, etc., and edges indicating the flow of information between these entities. Conversely, CTI reports sourced from various origins are transformed into query graphs, with nodes representing IoCs, files, processes, and more, and the relationships between nodes being derived from standard CTI-sharing sources. The graph alignment component employs an inexact graph pattern-matching algorithm to identify matches. Specifically, this component scrutinizes the presence of the query graph within the provenance graph. Poirot does not provide a comprehensive list of aligned paths; instead, it halts when the query identifies the first match that exceeds a predefined threshold. Additionally, constructing Poirot’s query graph is a complex task that demands the expertise of a human operator.

CTP-DHGL [254] leverages multi-source data and a heterogeneous graph to predict *evolutionary cyber threats* by forecasting links within the graph. The graph is structured to encompass nodes categorized into six types, with the potential for up to 14 relations between node pairs. CTP-DHGL predicts newly emerging links within the graph, which can signify the emergence of malware variants, shifts in attacker tactics, or the evolution of vulnerabilities. The proposed CTP-DHGL framework comprises five components: (I) Heterogeneous Graph Construction: This step involves defining nodes as cyber objects and specifying the relationships between these objects; (ii) Heterogeneous Graph Embedding: It utilizes a hierarchical attention mechanism to maintain representations at both the node-level and meta-path level. These embeddings capture the state of the graph at a specific timestamp; (iii) Transforming Embedding Graph into Temporal Sequence: The embeddings from the previous step are transformed into a temporal sequence; (iv) LSTM Encoder: it learns patterns of evolution within the graph based on the embedded sequence, and (v) Decoder: The decoder dynamically deciphers emerging patterns, revealing novel threat indicators. CTP-DHGL enhances threat prediction by applying a novel loss function that introduces a learnable penalty matrix. This ensures that the model effectively captures the dynamic evolutionary patterns in the data. In [255], Zhang et al. combined CTI with Security Requirements (SR) data to develop a knowledge graph and extracted new insights using a heterogeneous graph. They introduced EGNN, a model that generates independent edge representations and improves link prediction by propagating edge information, thereby enhancing the graph’s overall representation capability. The paper introduces the *evolving cybersecurity knowledge graph*, SEPSSES [256], that enables organizations to conduct vulnerability assessments effectively, helping them understand the impact of vulnerabilities on the organization and improving incident response by providing contextualized information. Specifically, this solution uses an Extraction, Transformation, and Loading (ETL) pipeline to integrate security data from various sources, both public and local, and to prepare and update a KG. SEPSSES consists of four key modules: (i) Data Acquisition Module: regularly polls diverse sources to ingest fresh information into the KG; (ii) Resource Transformation Component: developed in RDF Mapping Language (RML), converts data collected from different sources in diverse formats into a common ontology, (iii) Entity Linking and Validation Module: links data from different sources that carry identical identifiers. Furthermore, it validates the generated RDF using SHACL, and (iv) the Data Storage Module periodically executes a script to record the data in triple form, enhancing the knowledge generation process.

Predicting cyber *attack preferences* offers tremendous value, enabling security operators to formulate defensive strategies for mitigating attacks. The paper introduces a novel framework named HinAp [257],

which utilizes an attributed heterogeneous attention network and transductive learning. The primary goal of this investigation is to understand the characteristics of attackers, such as their intent to infiltrate systems, their use of intrusion tools, and their affiliation with malicious groups. Specifically, the authors create an Attribute Heterogeneous Information Network (AHIN) for an attack to model the attacker, vulnerabilities, exploits, compromised devices, and invaded platforms. They use BERT to convert the attributes of nodes (vulnerability nodes and attack script nodes) into 256-dimensional vectors. Afterward, they apply node-level attention within the AHIN to determine the weights of different nodes that characterize attack preferences. The authors then formulate 20 meta-paths and five types of meta-graphs to help explore the relationships between cyber objects. They then use transductive learning to build a meta-graph-based preference model. Finally, they stack these two models together to determine attack preference by investigating node-level and structural-level features comprising meta-paths and meta-graphs. APT-Scope [258] is a framework designed to predict APT groups by leveraging an enriched Heterogeneous Information Network (HIN) of CTI. It gathers APT-related data from OSINT sources and enhances it using methods such as DNS lookups, Whois queries, and NER with NLP models like SpaCy and CyNER. The enriched data is structured into a HIN, where nodes represent entities (e.g., APT groups, malware) and edges depict their relationships. Additionally, FastRP and Logistic Regression are employed to predict entity relationships and identify APT group aliases.

Piplai et al. [174] extract knowledge from malware After Action Reports (AARs) and transform them into a KG. Their proposed knowledge extraction pipeline comprises three components: Malware entity extraction, relationship extraction, and KG construction based on STIX. They integrate multiple AARs that refer to identical malware and similar attacks into the KG. Suppose the entities in the current document differ from those in previously processed AARs. In that case, the system calculates term similarity using the TF-IDF score for the current document and the AARs seen before. If a close match is found, the system fuses the graph created from the recent document with the existing CKG.

The author [259] proposes a hybrid model for *information extraction* by combining four important components: entity extraction, coreference resolution, relation extraction, and KG construction. To process a threat article, tokenization is performed, and BERT-based model encodings are obtained for each token. Additionally, POS embeddings are obtained for each token. POS-enhanced word embeddings are created by merging the word embeddings with the POS embeddings. For entity extraction, the POS-enhanced word embeddings are used as input for an attention layer to obtain contextual embeddings. Furthermore, embeddings are also generated for POS-enhanced word representations using Bi-LSTM. A linear classifier combines both the contextual embeddings and Bi-LSTM feature vector to produce entity labels. Coreference resolution is treated as a binary-class problem. POS-enhanced representations are obtained for each mention, and a CNN is employed to determine whether mentions refer to the same entity. Document-level RE is approached as a multilabel classification problem. This module utilizes a type embedding matrix and a distance embedding matrix to enhance entity relation extraction, resulting in an improved document representation. Finally, Neo4j is used to visualize relationships between entities mentioned in multiple documents. Experimental results indicate a 10.56% improvement in relationship extraction compared to state-of-the-art techniques.

More recently, in [260], the authors identify incorrect assertions (*outdated data/wrong information*) within the CKG by estimating scores for semantic triples (entity-head, relation, entity-tail). The conventional method for considering a triple as trustworthy involves recording temporal information about the collected data in the CKG. However, since time-related information is often absent in diverse CTI sources, the authors leverage GCN to process semantic triples without time-specific data. During the training phase of GCN, they adopt a supervised

approach where each semantic triple is associated with labels. The GCN then generates a probabilistic value between 0 and 1, indicating the relation between a pair of entities. A higher score is indicative of a more credible relationship. The proposed approach periodically updates these scores to filter out outdated triples or data originating from incorrect sources.

To address the issue of the potential inclusion of fake CTI within unstructured open-source information, which could lead to incorrect decisions and organizational harm, the authors in [261] have taken steps to enhance the existing CKG. This enhancement involves augmenting the CKG with *provenance graphs to ensure data authenticity*. Specifically, every node within the CKG is enriched with additional information that indicates its provenance. The process begins by creating a malware CKG, where entities related to malware are extracted using Conditional Random Fields (CRF) and regular expressions. Each entity within this graph is then subjected to a neural network classifier, which establishes links between each entity and one of six predefined relations. Next, the provenance system focuses on each node and constructs the provenance graph by augmenting it with associated information, which includes URL, Publisher, organization name, Author, Country, origin type, creation date, and provenance score. This comprehensive approach ensures the authenticity of the collected data. Subsequently, the provenance relation encoder is employed to derive relations within the provenance graph. Finally, the provenance fusion system merges the CKG with the provenance graph, thereby incorporating a trust dimension into the KG. This augmentation helps mitigate the risk of fake CTI and supports more reliable decision-making within the organization.

8.5. Explainability in CTI

Explainability plays a pivotal role in AI, serving as a vital bridge between the complexities of AI models and the understanding of end-users. While there is no universal definition for Explainable Artificial Intelligence (XAI), its core purpose remains consistent: to render AI outcomes comprehensible and transparent. In essence, XAI encompasses a range of methodologies aimed at empowering researchers to comprehend and trust ML outcomes [262]. It evaluates not only the accuracy of AI decisions but also their transparency and underlying rationale. This emphasis on understanding and trust is especially critical in the context of cybersecurity, where the adoption of AI has the potential to enhance defense mechanisms but must also maintain transparency to address evolving and intricate cyber threats effectively [263]. Incorporating XAI into cybersecurity models is an essential step toward achieving this balance, ensuring that AI-driven security systems are not only accurate but also comprehensible, enabling users to make informed decisions and effectively manage the complex landscape of cyber defense. In NLP, various XAI techniques are employed to enhance the transparency and interpretability of models. Local Interpretable Model-agnostic Explanations (LIME) generates localized explanations by perturbing input data, providing insights into why a model made a specific prediction. SHapley Additive exPlanations (SHAP) employs cooperative game theory to assign values to features, attributing importance to words or tokens in text. Attention mechanisms in transformer models like BERT and GPT reveal which words receive the model's focus. Counterfactual explanations generate alternative texts that would lead to a different model prediction, helping users understand how slight changes in input affect the output.

In [264], the authors employed explainable DL techniques to categorize onion services based on their content, including images and text. The core components of this approach consist of a CNN integrated with Gradient-weighted Class Activation Mapping (Grad-CAM) for image analysis and a pre-trained word embedding with Bahdanau additive attention for text analysis. They underscored the importance of explainability in facilitating well-informed decision-making and discussed misclassified samples and observations of false positive predictions.

Additionally, the paper highlights how certain false positives yield intriguing explanations, which were explained using Grad-CAM and attention visualization techniques. Li et al. [265] introduced a system that integrates threat intelligence and XAI to detect APTs. This intelligence not only offers valuable insights for refining defense strategies and resource allocation but also ensures the transparency and reliability of AI-driven detection outcomes. Their evaluation of the system involved the use of datasets, including KDD'99 and UNSW-NB15, and employed LIME for providing model explanations. In [266], authors presented an approach for APT attribution using paragraph vectors and bag-of-words vectors to represent function semantics and behavior reports, respectively. They used a Random Forest Classifier (RFC) and LIME to provide insights into the model's results.

XAI is also applied to explain the TTP classification task, as demonstrated in the work by Ge et al. [267]. Their method introduces a self-adversarial framework, consisting of an evidence generator and a TTPs classifier discriminator. This framework extracts evidence from CTI text through a topic prototype-based keyword importance filtering technique, resulting in both an evidence set and a perturbation set. A siamese discriminator is utilized to assess these sets' impact on TTP classification, ensuring that only elements from the evidence set are accurately categorized as TTP information. They compared various explanation methods and found that LIME's local semantic approach is less effective when contrasted with methods that consider global semantic contributions. Techniques relying on attention mechanisms or topic prototypes demonstrate enhanced prediction stability compared to SHAP. Furthermore, their approach, which integrates self-adversarial training and θ -sigmoid activation, proficiently eliminates irrelevant data and assigns greater significance to technical keywords. Within this range of methods, the Percent-M model, featuring a multi-topic prototype mechanism, showcases the capability to establish precise and unwavering decision boundaries while offering coherent textual explanations for classifying TTPs.

Lessons learned. From this section, we can observe that cyber entities are primarily extracted using a combination of models, such as BiLSTM, CRF, and BERT embedding. Researchers employ pipeline and joint entity recognition models to define relationships between these entities. The joint entity relationship techniques require further refinement. Visual representation of cyber security entities and their relationships is achieved through knowledge graphs. Researchers have employed tools such as Neo4j, Grakn, and Knowledge Graph Query Languages (Graql, SPARQL, Cypher) for this purpose. While XAI offers a valuable means to interpret model decisions, its widespread adoption remains limited. Incorporating XAI can significantly enhance the transparency and understanding of threat intelligence models

9. Cyber threat intelligence sharing and collaboration

This section explores the diverse landscape of CTI-sharing platforms, highlighting the various options available for organizations to collaborate and exchange critical threat information. Furthermore, we delve into the adoption of standards within CTI sharing, underlining the industry's concerted efforts to establish uniform practices for enhanced information sharing. Moreover, our exploration encompasses research works aimed at assessing the quality of CTI sharing and exploring the dimensions proposed to evaluate the effectiveness and reliability of shared intelligence data. Additionally, we reviewed the strategies and mechanisms employed to achieve privacy in CTI sharing, emphasizing the importance of safeguarding sensitive information while fostering collaborative efforts to combat cyber threats.

9.1. CTI information-sharing platforms

Traditionally, organizations relied on informal means like phone calls and emails for sharing threat intelligence. However, there is a

growing trend toward establishing connected communities that employ dedicated platforms to facilitate the automated or semi-automated sharing of threat intelligence. Threat Intelligence Platforms (TIPs) are specialized software that assists organizations by gathering, correlating, and analyzing real-time threat information from various sources to strengthen their defensive strategies. While numerous TIPs are available in the market, most are offered under commercial licenses. Virus-Total is the leading intelligence service, employing over 80 antivirus engines, sandbox systems, and blacklists to analyze URLs and files for a wide range of malicious content and produce threat reports. MISP [268], known as the Malware Information Sharing Platform, functions as an open-source solution created to gather, store, distribute, and exchange cybersecurity indicators and threat data. It streamlines the efficient exchange of information both within the security community and beyond, offering a range of features, including an indicator database, automated correlation, adaptable data modeling, sharing capabilities, user-friendly interfaces, and compatibility with various data formats and standards. In [269], Stojkovski et al. evaluated security information sharing on the MISP platform and the importance of considering user experience to meet the varied needs of security information professionals. The investigation employed a combination of methodologies, such as the user experience questionnaire and sentence completion, to assess MISP's advantages and drawbacks. The findings indicate that the platform leans toward a technical orientation with a steep learning curve, which may pose challenges for newcomers. In two recent works [270–272], the authors leverage MISP for efficiently share CTI IOCs. In particular, [270] propose an extension of the MISP platform called CLEVER that incorporates three ML-based modules, namely (i) threat ranking, (ii) event classification, and (iii) IoC clustering. Additionally, a policy-based method for assigning levels of confidentiality to IOCs is introduced, to assess data privacy and secure sharing in collaborative scenarios. The authors of [271], instead, first create detailed attack information in the form of IoC indicators when a security incident involving IoT devices occurs, then by applying a "Hub & Spoke" model within the MISP system, the event information is shared, enabling efficient responses to potential threats that may arise in the future. Finally, the system proposed by Valdés et al. [272] aims at enhancing the MISP platform by integrating a risk assessment tool, which enriches the shared intelligence with detailed risk information, supporting informed decision-making.

In another study [273], Jollès et al. explored three collective intelligence dynamics on the ThreatFox platform, emphasizing the significance of trust networks in participant onboarding, highlighting the platform's growth patterns in IoC publication, and discussing the potential of a credit system as an incentive for information sharing in the cybersecurity community. ThreatFox, managed by abuse.ch, is a unique open-data TIP that promotes IoC sharing through accessibility, user-friendliness, and a public good approach. Similar to ThreatFox, abuse.ch has created additional initiatives focused on monitoring and disseminating information regarding cybersecurity threats. These projects encompass URLHaus, designed for the exchange of malicious URLs, MalwareBazaar, which collects and distributes malware samples, and YARAify, a tool for scanning suspicious files using YARA rules. Sauerwein et al. [18] assessed nine TIPs such as ThreatStream, ThreatQ, ThreatConnect, Open Threat Exchange (OTX), MISP, IBM X-Force Exchange, Falcon X Intelligence CrowdStrike, Collective Intelligence Framework (CIF), and Collaborative Research into Threats (CRITs) by examining how they align with the CTI life cycle. Their investigative case studies uncovered that the current focus of these platforms is primarily on the pre-processing and dissemination stages. A similar study by De et al. [274] assessed Anomali STAXX, OpenCTI, CRITs, CIF, and MISP based on architectural and CTI processes. Their findings indicated that among these TIPs, MISP and OpenCTI were regarded as the most comprehensive and adaptable solutions. OpenCTI is an open-source platform that adheres to STIX2

standards, providing a modern Web application with integration capabilities, enabling the seamless connection and processing of both technical and non-technical threat information. Additionally, it supports data export in various formats, including CSV and STIX2 bundles. The authors also highlighted that while advanced solutions exist for CTI, none of them cover the complete CTI process. Despite the presence of outstanding options, the search for a comprehensive solution for threat intelligence-based defense remains challenging due to the diverse focuses of these platforms, which address only specific stages. Wagner et al. [275] analyzed 30 TIPs; they found that AlienVault offers anonymity but lacks content masking, HP Threat Central allows preprocessing and content masking, and Comilion provides architecture-level anonymity but could be susceptible to Personally Identifiable Information (PII) analysis. Their approach involves storing CTI locally, utilizing Java regular expressions and exceptions for anonymization, and enabling sharing through an MISP instance with TOR network anonymization, permitting customization following the Traffic Light Protocol (TLP). However, the evaluation revealed vulnerabilities to background knowledge attacks.

Haque et al. [276] emphasized the importance of adopting a controlled and automated strategy for sharing CTI while highlighting the effectiveness of Relationship-Based Access Control for facilitating this sharing. They introduced an automated method for identifying, generating, and disseminating structured CTI and demonstrated these concepts through a prototype Automated Cyber Defense System in a cloud-based environment. ETIP [277], a platform designed for the purpose, is geared toward the collection and processing of structured data sourced from external sources, encompassing OSINT feeds, along with information originating from an organization's network infrastructure. ETIP comprises three core modules: the input module, responsible for the collection and standardization of IoCs from OSINT feeds and the monitoring infrastructure; the operational module, which correlates the data to generate enriched IoCs and evaluates threat data using a threat score; and the output module, tasked with visually presenting outcomes and sharing them with external entities to bolster cybersecurity defenses. IoC normalization and representation in MISP format are facilitated through the use of the MISP. The platform employs a deduplication technique employing *contained similarity* to eliminate duplicate IoCs, consolidates related IoCs into composed IoCs and assigns a threat score to each IoC. This threat score enables Security Operations Center (SOC) analysts to prioritize security incident investigations based on the threat score value, which factors in heuristic weights for individual attributes and overall complexity.

Sharing threat intelligence has the potential to improve IT security, but it faces obstacles such as expenses, risks, and legal requirements. To overcome these challenges, Menges et al. [278] presented DEALER, a platform that promotes secure threat intelligence sharing by providing incentives and addressing compliance concerns. DEALER, which operates on the EOS Blockchain and InterPlanetary File System (IPFS) distributed hash table, facilitates the anonymous sharing of well-structured incident data. It utilizes impartial quality metrics for reputation assessment and offers safeguards for both buyers and sellers through dispute resolution and cryptocurrency rewards. However, the authors recommended limiting the platform's use to sharing noncritical data, as it may not be suitable for sharing highly sensitive and critical CTI. In response to the challenges related to trust in the source and integrity of threat intelligence data, Preuveneers et al. [279] improved the security framework TATIS [280]. This framework guarantees that only individuals with authorization can access sensitive data when it is being transferred between various threat intelligence systems. The practical implementation of this approach involves a distributed system built on top of the MISP, where distributed ledger technology (DLT) is employed to manage access to CTI data. To safeguard the shared data, encryption is applied using the Ciphertext-Policy Attribute-Based Encryption (CP-ABE) cryptographic scheme.

9.2. Standardization efforts

The structured representation of CTI data is beneficial, rather than sharing simple plain text. The primary aim of CTI standards is to provide a thorough methodology for describing the details of threats, attacks, and security incidents. These standards can also have specialized purposes beyond just representing security information. They may be designed for specific operational scenarios in the cybersecurity field. One widely recognized standard for sharing CTI data is Structured Threat Information Expression (STIX), launched by the US government and MITRE [281]. STIX acts as a format for data exchange rather than a model for data storage. This standardization, developed in collaboration with a wide range of stakeholders, specifies, captures, characterizes, and communicates standardized information regarding cyber threats. It fulfills various use cases, including analyzing cyber threats, defining indicator patterns for cyber threats, managing cyber threat response activities, and sharing CTI within communities in a standardized format. STIX utilizes JSON format to represent CTI objects and relationships. STIX 2.1 comprises 18 objects such as Attack Pattern, Campaign, Course of Action, Grouping, Identity, Indicator, Infrastructure, Intrusion Set, Location, Malware, Malware Analysis, Note, Observed Data, Opinion, Report, Threat Actor, Tool, Vulnerability. In [45,72,174,178,181,216,236,239], STIX is used as the standard for defining entities and relationships in entity-relationship tasks. Fig. 11 illustrates the representation of threat information in the STIX format, sourced from the STIX Web site.⁷ In this example, there are two distinct objects: the first is a *threat actor* for "Disco Team Threat Actor Group", and the second is an *identity* type for "Disco Team". These two objects are interconnected through an *attributed_to* relationship. Also, the illustration provides additional information, including alias names, contact details, a description, and more, to provide a comprehensive overview of these objects. In [282], Riesco et al. incorporated an integrated framework for conducting Dynamic Risk Assessment and Management (DRA/DRM), employing the Web Ontology Language (OWL), a semantic reasoner, and the Semantic Web Rule Language (SWRL). They developed a semantic adaptation of STIX™v2.0 by integrating OWL and SWRL, facilitating the automatic generation of dependencies and the detection of threats. Data formatted in STIX can be transmitted using various communication protocols. One such protocol is Trusted Automated eXchange of Indicator Information (TAXII), which operates at the application layer and is specifically designed to exchange threat information structured in the STIX format. TAXII, initiated by MITRE, supports various sharing models, encompassing publish-subscribe, peer-to-peer (P2P), and hub-and-spoke approaches. Another structured language, Cyber Observable eXpression (CybOX), was developed for representing cyber observable events. With the introduction of STIX 2.1, CybOX has been seamlessly integrated and is now an integral part of the STIX standard. Furthermore, MITRE provides MAEC (Malware Attribute Enumeration and Characterization), a highly versatile language for sharing information about malware. MAEC enables the encoding and transmission of comprehensive information regarding malware, including its characteristics, such as attack patterns, artifacts, and behaviors.

VERIS serves as a standardized metrics framework aimed at establishing a universal language for systematically explaining security incidents. It tackles a prominent issue in the security sector (the scarcity of high-quality information) by aiding organizations in the collection of valuable incident-related data and its responsible and anonymous sharing with others. VERIS is organized around four fundamental components: Actor, Action, Asset, and Attribute, with a primary emphasis on documenting and evaluating internal incidents. It includes incident

⁷ <https://oasis-open.github.io/cti-documentation/examples/identifying-a-threat-actor-profile>

```

{
  "type": "bundle",
  "id": "bundle-601cee35-6b16-4e68-a3e7-9ec7d755b4c3",
  "objects": [
    {
      "type": "threat-actor",
      "spec_version": "2.1",
      "id": "threat-actor-dfaa8d77-07e2-4e28-b2c8-92e9f7b04428",
      "created": "2014-11-19T23:39:03.893Z",
      "modified": "2014-11-19T23:39:03.893Z",
      "name": "Disco Team Threat Actor Group",
      "description": "This organized threat actor group operates
        to create profit from all types of crime.",
      "threat_actor_types": [
        "crime-syndicate"
      ],
      "aliases": [
        "Equipo del Discoteca"
      ],
      "roles": [
        "agent"
      ],
      "goals": [
        "Steal Credit Card Information"
      ],
      "sophistication": "expert",
      "resource_level": "organization",
      "primary_motivation": "personal-gain"
    },
    {
      "type": "identity",
      "spec_version": "2.1",
      "id": "identity-733c5838-34d9-4fbf-949c-62aba761184c",
      "created": "2016-08-23T18:05:49.307Z",
      "modified": "2016-08-23T18:05:49.307Z",
      "name": "Disco Team",
      "description": "Disco Team is the name of an organized
        threat actor crime-syndicate.",
      "identity_class": "organization",
      "contact_information": "disco-team@stealthemail.com"
    },
    {
      "type": "relationship",
      "spec_version": "2.1",
      "id": "relationship-a2e3efb5-351d-4d46-97a0-6897ee7c77a0",
      "created": "2020-02-29T18:01:28.577Z",
      "modified": "2020-02-29T18:01:28.577Z",
      "relationship_type": "attributed-to",
      "source_ref": "threat-actor-dfaa8d77-07e2-4e28-b2c8-92e9f7b04428",
      "target_ref": "identity-733c5838-34d9-4fbf-949c-62aba761184c"
    }
  ]
}

```

Fig. 11. STIX 2.1 example representing entity relationship.

narratives, local incident identifiers, data on affected users, and response actions taken. Currently, VERIS is in version 1.3.1, utilizing the JSON format. Incident Object Description Exchange Format (IODEF) serves as a structured data framework created to facilitate the exchange of information concerning computer security incidents. It is frequently employed by Computer Security Incident Response Teams (CSIRTs) and includes both an information model and a corresponding data model specified through XML Schema. Common Vulnerability Reporting Format (CVRF) stands out as another standard in cybersecurity, designed with a machine-readable format to facilitate structured creation, updating, and exchange of security advisories. It encompasses data related to products, vulnerabilities, and the current state of impact and remediation, thereby promoting efficient communication within the cybersecurity community. The embrace of CVRF by MITRE's CVE repository, the central hub for vulnerabilities and exposures. CVE is designed to recognize and label vulnerabilities in hardware and software products, with each CVE entry including a distinct identifier, a vulnerability description, and links to supplementary information. Open Vulnerability and Assessment Language (OVAL) is a globally recognized initiative, freely available to the public, aimed at standardizing the assessment and reporting of computer system states in the field of information security. Its primary purpose is to provide a language and framework for standardizing the assessment and reporting of computer system states, with a particular emphasis on system vulnerabilities and configuration concerns. While OVAL is not explicitly classified as a CTI standard, it can be effectively combined with CTI standards and tools to bolster cybersecurity practices and elevate the sophistication of threat analysis.

The study [19] emphasizes the significance of a systematic approach and cultivating shared understanding in the development of incident response standards. Also observing a growing shift toward structured

incident response formats due to the increasing prevalence of Security Orchestration, Automation, and Response (SOAR) products. In [283], Menges et al. assessed and compared various incident reporting formats, which included STIX, X-ARF, VERIS, and IODEF, with different versions. Their primary focus lies in evaluating general criteria such as aggregability, extensibility, interoperability, and readability and additional aspects like licensing and documentation costs. Also, this paper provides a more detailed examination of structural evaluation criteria, encompassing indicators, attackers, attacks, and defenders, and illustrates how these criteria relate to typical security use cases. Remarkably, their findings identified STIX as the most comprehensive and practical choice for reporting cybersecurity incidents. The study [284] evaluated specific CTI standards like MISP and STIX; they also included common formats such as RSS and CSV in their examination. The research revealed that even though there are established standards available, many key supporters and CTI producers still opt for creating their own simplified JSON formats, regardless of the time and formal processes required for standardization. Further, the study suggested that for most security use cases, a recommended approach might rely on the IDEA format, supplemented by MISP or STIX, as it proved effective in meeting CTI requirements. The study [285] examined distinct attributes of different sharing standards, such as STIX, IDEA, VERIS, IODEF, OpenIOC, and MISP's Format. They evaluated three categories of risks related to attribute disclosure: direct disclosure, indirect or inference risk, and accidental disclosure. Direct disclosure pertains to situations where an attribute immediately reveals sensitive information, while indirect risk comes into play under specific conditions, like having a sufficient number of related events or being susceptible to reidentification attacks. Accidental disclosure is a concern for attributes situated in fields with less stringent formatting. Excluding VERIS, the analysis showed that the risk of attribute disclosure had an upper limit of approximately 20%, with direct disclosure typically being below 10%, especially in the case of the two most widely used formats, STIX and MISP. VERIS displayed a high risk of accidental disclosure due to its emphasis on detailed documentation, whereas IODEF presented lower risks due to its structured approach. STIX carried higher risks of accidental disclosure or inference due to its uncontrolled input, while OpenIOC benefited from well-defined fields, mitigating potential accidental disclosure. MISP utilized secure identifiers but featured free-text fields with accidental disclosure concerns.

9.3. Quality and privacy in CTI sharing

Evaluating the quality of CTI is essential for its effective sharing and utilization, as it directly impacts incident response times. Research findings emphasize the need to maintain CTI quality during collaboration [286]. They have pointed out that assessing the quality of CTI poses a major challenge in contemporary cybersecurity information-sharing contexts. When considering integrating a CTI source into an organization's security tools, one must consider various attributes, including the quality of the data or information it provides. In [287], Geras et al. interviewed experts to examine how practitioners handle CTI quality, emphasizing that high-quality CTI is crucial for making informed decisions. The required quality level depends on the type and application of CTI, and context information is key to its effective use. In [288], the authors proposed a relevant set of quality dimensions and established corresponding metrics for evaluating CTI. They also differentiated which metric could be computed automatically and which needed security experts to assess. They categorized the quality dimensions into three levels: the Attribute Level, which focuses on specific attributes of STIX objects (e.g., evaluating timeliness using attributes like modified and created); the Object Level, which is not restricted to predefined attributes and can be assessed based on variable attribute sets (Representational Consistency) or the entire object itself (Reputation); and finally, the Report Level, which addresses experts' need to determine if a report contains an appropriate

amount of data. Also, Zibak et al. [289] identified different quality dimensions: *timeliness*, *reliability*, *relevance*, *provenance*, *interoperability*, *actionability*, and *accuracy* through a systematic review of existing CTI literature, subsequently refining and validating them via a modified *delphi* study. In another study, Qiang et al. [290] conducted a multi-dimensional quantitative evaluation of threat intelligence services with a user-centric perspective. Experts assess the criteria manually through a questionnaire and adjust their significance using a multi-objective algorithm to derive final ratings. Although this method primarily focuses on commercial CTI providers, it does not explore the granular aspects of threat intelligence. Empirical findings suggest that this approach excels in evaluating the comprehensiveness and depth of threat intelligence data feeds as well as providing an overall situational assessment.

In [291], authors assessed the credibility of CTI sources through quantitative parameters and employed a customizable weighted evaluation approach to enhance trust in these sources. The mechanism operates within a closed world assumption, evaluating each source's trustworthiness relative to others using parameters such as *similarity*, *completeness*, *timeliness*, *compliance*, *interoperability*, *intelligence*, *verifiability*, *positives*, *false*, *maintenance*, and *extensiveness*. Their approach can dynamically reassess parameters when new threat intelligence is shared, allowing for real-time trust adjustments without needing expert intervention in source validation. Griffioen et al. [292] evaluated the quality of 17 open-source CTI feeds over 14 months and seven more feeds over seven months. They evaluated based on four quality dimensions such as *asoriginality*, *impact*, *sensitivity*, and *timeliness*. The findings reveal significant variations in performance among these feeds, providing timely and effective indicators while others exhibit delays in reporting and minimal influence on malicious activities. In [293], Chen et al. evaluated the quality of IoCs by using a STIX graph-based methodology, combining various threat intelligence sources to enhance the precision and relevance of threat severity and confidence scores. They utilized a heuristic evaluation framework that assigns weighted values across five essential factors such as consistency, completeness, timeliness, relevance, and accuracy. Furthermore, they incorporated models like BERT-CRF for scoring relevance and a modified PageRank algorithm for ranking sources. Jin et al. [294] developed CTI-Lense to assess the effectiveness of CTI sharing using STIX data. Over nine years (2014–2023), they collected 6 million STIX objects from 10 sources and evaluated metrics such as quality, coverage, timeliness, and volume. The study revealed that despite increased data sharing, over 37% of the data within sources was duplicated. While URL data was shared promptly, facilitating rapid threat detection, malware signatures experienced significant delays. Furthermore, the analysis uncovered that 19% of threat actor data is incorrect.

AttTucker [295], a graph quality assessment model based on transformers, employs multiple attention heads to capture information about entities and their relations in a low-dimensional space. In the encoding phase, the model constructs query, key, and value vectors using the embeddings of entity heads and relations, with the aid of multiple attention heads, to generate the output representation. In the decoding phase, it calculates the predicted probability for each triple. The probability vector is then concatenated with path-level information and fed into a multi-level perceptron to derive the final score for the triple. In [296], the authors assessed the quality and availability of CTI sources by employing a correlation graph and an iterative algorithm to evaluate feed trustworthiness. Their proposed framework features a CTI Generator for structuring intelligence, a Feed Assessment to evaluate trustworthiness, and a Content Assessment to rate quality based on criteria such as verification, timeliness, and completeness, ultimately culminating in a Comprehensive Assessment model that adjusts to network dynamics. While the framework offers a robust approach, it lacks considerations like coverage and interpretability, and further research is needed to incorporate user feedback and manual evaluation to improve credibility assessments.

The significance of trust in threat information sharing arises from the presence of private and sensitive data in CTI. Many organizations are cautious about sharing or obtaining cybersecurity threat information because they fear becoming targets for threat actors. Threat can lead to damage to their reputation, regulatory investigations, and legal conflicts. In [297], the authors introduced a trust system designed to establish a secure environment for sharing threat intelligence. They thoroughly examined popular CTI platforms and providers and validated the effectiveness of their trust framework through practical case studies. Chadwick et al. [298] proposed a trust model comprising five distinct levels within a cloud-edge data-sharing infrastructure to secure the organization's CTI data. At Level 1, organizations exhibit complete trust in all involved parties and share their CTI data as-is in plain text. Level 2 extends trust primarily to the cloud infrastructure provider but harbors partial trust in collaborators, permitting the provider to apply anonymization or pseudonymization to CTI data. Level 3 maintains full trust in the provider while assuming untrusted collaborators, necessitating the encryption of CTI data by the provider, specifically homomorphic encryption for analysis. Level 4 signifies partial trust in all parties, allowing data sharing with the condition that sensitive fields are anonymized or pseudonymized beforehand. Lastly, at Level 5, organizations express no trust in collaborators or the cloud provider conducting the analysis, mandating the homomorphic encryption of CTI data before sharing.

Several researchers [299–303] have leveraged Blockchain technologies to achieve privacy in the sharing of CTI. Homan et al. [300] presented a CTI sharing model using Hyperledger Fabric Blockchain, addressing trust issues by trusted communities through Blockchain channels and enforcing Traffic Light Protocol (TLP) sharing rules via smart contracts within the network. DefenseChain [301] employed Blockchain to enhance trust in CTI sharing among various organizations. The implementation of DefenseChain took place on an Open Cloud testbed, using Hyperledger Composer, as well as in a simulated environment. Riesco et al. [302] employed the Ethereum Blockchain, along with the standard CTI token, to promote the exchange and dynamic sharing of threat intelligence among various stakeholders. They integrated semantic Web standards to streamline the sharing of knowledge related to behavioral threat intelligence patterns. Huff et al. [303] has devised a solution that allows for secure information sharing and enables entities to participate in threat-sharing communities anonymously using a distributed Blockchain ledger. The incorporation of zk-SNARKs and Sparse Merkle Trees further strengthens anonymity and the efficient utilization of tokens. BFLS [299] has integrated Federated Learning (FL) [304] with Blockchain technology to create a decentralized approach for training a threat detection model. They enhance the Blockchain's consensus protocol to select high-quality CTI models for participation in FL. The chosen models are automatically aggregated and updated using smart contracts, improving model quality and more efficient sharing. In [305], authors developed SeCTIS, a system that combines Swarm Learning and Blockchain to enable businesses to work together securely while keeping their CTI data private. They also developed a way to assess both the quality of data and models and how trustworthy the participating organizations are, using validators with Zero-Knowledge Proofs. In [306], Rahman et al. utilized encrypted FL to enable secure collaboration among multiple organizations while protecting sensitive data. Their system also incorporated a malicious client detection mechanism that continuously evaluates local model performance to identify and counteract man-in-the-middle attacks and data poisoning attacks.

Dunnett et al. [307] proposed the Priv-Share framework, which combines differential sharing and trustless delegation to improve the privacy and efficiency of CTI sharing. The approach involves dividing sensitive CTI into smaller data groups that can be shared selectively based on trust metrics, giving organizations control over the disclosure of sensitive information. It also features a decentralized trust evaluation

process managed by democratic group leaders and includes an incentive mechanism to enable CTI producers to delegate trust assessments without financial risk. The framework was validated through analytical methods using queuing and game theory, alongside a proof-of-concept implementation on the Ethereum blockchain.

Lessons learned. From this section, we can observe that the establishment of standardized quality metrics and the enhancement of data-sharing privacy are imperative for building trust and promoting collaboration within the cybersecurity community. Additionally, extracted TTPs and threat intelligence are often shared in standardized formats, like STIX, to enhance collaboration among cybersecurity professionals. Moreover, the integration of Blockchain and Federated Learning approaches can significantly enhance privacy.

10. Adversarial attack

CTI is essential for timely threat warnings, but if adversaries compromise it, it loses its effectiveness. CTI is typically derived from publicly available sources such as security blogs, Twitter hacker forums, the Dark Web, and security bulletins. These platforms often provide advance notice of vulnerabilities, which can be compared to standard sources like NVD or CVE. Unfortunately, adversaries can introduce counterfeit threat data to disrupt the creation of actionable CTI. This misinformation can cause the AI cyber defense system to generate contradictory information, leading to incorrect policy updates. This, in turn, can harm network and endpoint defensive rules, leading to the selection of diversionary actions to evade attacks, ultimately putting the organization at risk. Adversarial ML techniques have become a preferred tool for malicious insiders. These techniques feed deceptive inputs to ML models, causing them to make incorrect decisions or predictions. Just as adversarial attacks have gained prominence in image recognition systems, much research focuses on malicious attacks on text data.

10.1. Adversarial attacks in NLP

In [308], Papernot et al. conducted the initial study on adversarial attacks targeting text. To perform the attack, they embraced the strategy of crafting adversarial images and produced adversarial text samples by manipulating the derivatives linked to text embeddings. Earlier, researchers adapted image-based adversarial techniques to maintain semantic and syntactic integrity in a text through methods like the Greedy Search algorithm [160] and Reinforcement Learning [309]. Later, they developed specialized adversarial attacks for text due to their discrete nature. Creating adversarial samples involved changing individual characters and words, effectively preserving semantic consistency and syntactic correctness. However, this approach often resulted in a lack of diversity in the generated adversarial texts. Later, researchers shifted their focus to modifying entire texts to achieve greater diversity while maintaining semantic and syntactic integrity. These methods, in general, require careful design to ensure both effective attacks and high-quality adversarial texts. In the remainder of this section, we will describe recent approaches dealing with adversarial attacks in NLP, examples generation, poisoning attacks, and defense techniques. Observe that, we analyze all the available papers in the literature according to our selection criteria, and the small number of identified works comes from the effort distribution in the literature.

10.2. Adversarial examples generation

An adversarial example is a deceptive text resembling a benign input designed to deceive AI models. Compared to the generation of adversarial images, crafting adversarial text is more challenging. Various strategies are used in constructing adversarial instances for NLP tasks, encompassing alterations at the character, word, and sentence

<p>ORIGINAL TEXT</p> <p>Tactic: Reconnaissance</p> <p>Technique: Spearphishing Attachment</p> <hr/> <p>Adversaries may send spearphishing messages with a malicious attachment to elicit sensitive information that can be used during targeting.</p>
<p>ADVERSARIAL TEXT</p> <p>Tactic: Reconnaissance</p> <p>Technique: Spearphishing Link</p> <hr/> <p>Adversaries may send spearphishing messages with a malicious link to elicit sensitive information that can be used during targeting.</p>

Fig. 12. Adversarial example of a CTI text.

levels [310,311]. Common strategies entail character replacement, the addition of random characters or punctuation marks, the removal of characters or punctuation, character exchange, and the random reshuffling of character order within tokens. In image perturbation, even subtle pixel changes can render the image imperceptible to the human eye. In contrast, text-based adversarial examples can often appear visually or semantically similar to genuine content. When we apply character-level perturbation, these perturbations become more noticeable and easier to detect [312]. For instance, spell-checker tools can easily identify typographical errors, and if words or characters are altered, it can result in grammatical errors or change the meaning that can be detected. Aside from character-level changes, several strategies focus on modifying the text's semantic content. This involves substituting synonyms or antonyms for verbs, adverbs, and adjectives, as well as altering verb negations or tenses. Remarkably, even the elimination of frequently used stop-words can impact the performance of models. Another strategy entails injecting various sentence types into the text, such as grammatically correct human-approved or ungrammatical sentences. Adversarial attacks may also include rephrasing statements or intentionally removing essential sentence components. These strategies aim to provide small but effective changes to text inputs to mislead DNN models in NLP tasks. The Fig. 12 illustrates an adversarial example generated from the CTI text describing an attack technique, specifically *spearphishing attachment*. A slight modification has been made to create this adversarial example, primarily involving the substitution of "malicious attachment" with "malicious link". This substitution maintains the overall meaning of the sentence, but it transforms the adversarial example to describe a different attack technique, namely, *spearphishing link*. This change in wording can influence how individuals perceive and respond to the threat, emphasizing the importance of being cautious of both email attachments and suspicious links. Such adversarial attacks can impact tasks like attack tactics and technique classification.

10.3. Attacks and defense techniques in CTI

Adversarial techniques encompass various categories, including poisoning attacks, inversion, functional extraction, and evasion. Among these, data poisoning attacks are notably prevalent and disruptive in the context of CTI. One specific example of such an attack is the VirusTotal poisoning attack, exemplified by the McAfee Advanced Threat Research team [313]. Since many intrusion detection and security research efforts rely on data from VirusTotal, such attacks can significantly impact the CTI process. In a data poisoning attack, the objective is to contaminate the training data of AI systems through various methods. This malicious contamination can involve the inclusion of fake data into the AI system's training corpus. By doing so, attackers can manipulate the system's learning process, causing it to perform poorly on genuine data. For instance, in a recent study [78], researchers utilized the GPT-2 model to generate fake CTI data. They collected data from security blogs, vulnerability databases, and APT reports. This collected data was used to fine-tune the GPT-2 model, which was then employed to produce fabricated CTI data. The generation process was initiated with

a provided prompt, which served as input to the fine-tuned GPT-2. This model utilized tokenized prompts, normalization layers, attention layers, feed-forward neural networks, and softmax layers to produce fake CTI data. The generated fake CTI data was subsequently utilized to demonstrate a data poisoning attack on a knowledge extraction system.

In [314], the authors introduced a word-substitution-based adversarial attack targeting CTI NER systems. They first identified keywords for modification using POS tagging with the Flair NLP library. Then, synonyms for substitution are retrieved from WordNet and OpenHowNet, with a cosine similarity threshold of 0.2 applied to ensure semantic relevance in the cybersecurity context. To preserve grammatical integrity, the Pattern library was employed to adjust inflections such as tense and plurality according to POS tags. Randomized word selection further increased the attack's unpredictability, generating diverse adversarial examples. The authors analyzed the transferability to determine how well adversarial examples designed for one model could deceive others. The authors also compared their approach against existing adversarial text attack methods, including DeepWordBug [315], Morpheus [316], and TextFooler [317]. Their method achieved an attack success rate with an average transferability of over 20.83% across multiple NER models. To counter the word-substitution-based adversarial attack, they implemented a defense mechanism using Semantic Role Labeling (SRL). The SRL model processes each sentence by assigning semantic roles to its components and analyzing their dependencies. By examining these relationships, the model detects disruptions in expected role structures caused by adversarial modifications. Sentences exhibiting inconsistencies or deviations from the expected semantic patterns are flagged as potential adversarial texts. This approach achieved an average accuracy of 60.12% in detecting adversarial examples across three NER models.

In [318], Mongardini et al. simulated adversarial attacks targeting Intellectual Property (IP) theft. The authors assessed the effectiveness of their DARD system in misleading automated analysis tools used by attackers for classifying and extracting stolen data. Security practitioners, acting as ethical adversaries, conducted these attacks in a controlled environment. The process involved preprocessing exfiltrated documents through tokenization, stemming, and TF-IDF feature extraction to convert text into structured data. The attackers then applied K-means clustering to group similar documents and used LDA to identify key topics within clusters. Ideally, clustering would accurately group related documents, and topic modeling would extract meaningful keywords, enabling attackers to locate high-value IPs. However, DARD disrupts this by introducing deceptive operations, such as Basic Shuffle, Shuffle Increment, Shuffle Reduction, and Change Topic, which replace key terms with misleading keywords and alter document relationships. This results in incorrect clustering and topic modeling. Experiments on real-world datasets, like ArXiv papers, show that different types of attackers such as Black Box (naïve), Gray Box (partially aware), and Enhanced Gray Box (using an Oracle function) struggle to recover the original document structure. Topic modeling attacks using LDA and tools like Amazon Comprehend only retrieve deceptive keywords. They used the Adjusted Rand Index (ARI) to evaluate the attack. The adversaries achieved low clustering accuracy, demonstrating that DARD forces attackers to rely on manual inspection rather than automated classification.

In [319], the authors implemented a data poisoning attack to generate fake CTI. They employed a generative adversarial network, FTIGAN, to create 13,000 fabricated CTI data. Data was gathered from security blogs, APT reports, CVE vulnerability databases, and open-source threat intelligence sources, then structured in the STIX format and processed using CTIToken for extracting key cybersecurity entities like attack techniques, malware names, and threat actor groups. The FTIGAN model consisted of a generator, which used BERT-based word embeddings, an LSTM layer, an attention mechanism for contextual accuracy, and a discriminator built with TextCNN to assess the authenticity of the generated CTI. These fake CTIs were combined with 4,000 real

CTIs to form a composite dataset, which was used to train a model aimed at detecting fake CTI. The attack's success was evident in the alterations to knowledge graphs, where relationships between threat actors, attack techniques, and targets were significantly disrupted. To detect the fake CTI, the authors developed a Fake CTI Detector utilizing a transformer architecture, with features such as positional encoding, attention mechanisms, and feedforward neural networks.

To tackle the adversarial attack, [320] proposed a method that employs a Support Vector Classifier (SVC) and an embedding model. This method assesses the reputation of threat intelligence before integrating it into the AI system. Specifically, this work aims to determine the credibility of Reddit posts and 'Redditor features'. The authors collected a diverse corpus of cybersecurity data and converted each Reddit post into a vector using a 'bag of words' representation. The embeddings for each word were obtained from an existing system known as 'Cyber-All-Intel'. Consequently, each post is represented as the sum of cybersecurity term vectors. From a large collection of posts, a ground truth dataset was developed through manual labeling. The authors created two clusters: one consisting of 'credible' posts and the other referred to as the 'non-credible' set. They developed a linear SVC model, which was trained on 1206 posts using both post and Redditor features. Finally, the ensembles of the SVC model and the word embedding model were combined to compute the reputation score and validity of the posts. The proposed system achieved an accuracy rate of 71.73% in classifying posts as 'credible' or 'non-credible'.

Lessons learned: In conclusion, we can assess that, as the primary source of CTI data is social media and security blogs, it is inherently susceptible to adversarial attacks. Developing efficient adversarial defense techniques is a critical need in this field

11. Role of large language models in CTI

Large language models (LLMs) excel in understanding complex contexts and reasoning. In [321] authors leverage ChatGPT, to construct knowledge graphs from Cyber Threat Intelligence (CTI) by automating the extraction of attack-related entities and relationships from 13 real-world reports. A CTI-specific ontology, based on STIX v2.1, ensures structured results, guiding the LLM in metadata extraction and SPO triplet generation. ChatGPT's zero-shot information extraction capabilities demonstrate its effectiveness in automating structured CTI knowledge generation. LLM-TIKG [322] employs an LLM to construct a knowledge graph from 12,545 unstructured open-source threat intelligence reports. Leveraging GPT's few-shot learning capabilities, data is annotated and augmented (through translation and rephrasing) to create datasets for fine-tuning a smaller LLM (Llama 2-7B). Specifically, the 'GPT-3.5-turbo' model is utilized with an "instruction + examples + input" prompt template to generate the desired output structure. The fine-tuned model is then used to extract topic classifications, entities, and relationships, and ultimately, derive Tactics, Techniques, and Procedures (TTPs) from attack descriptions. The study in [323] focuses on automating the extraction of actionable CTI using LLMs and KGs. Open-source LLMs such as Llama 2, Mistral 7B Instruct, and Zephyr were evaluated with techniques such as few-shot learning, the guidance framework, and fine-tuning. The best-performing model, identified through ROUGE scores and human evaluation, was used to construct KGs for link prediction tasks, enabling proactive threat detection by forecasting future cyber threats, such as novel malware attack patterns. The guidance framework outperformed prompt engineering with higher ROUGE scores and ease of implementation. The study highlighted the need for better prompts, larger datasets, and improved link prediction for enhanced performance.

Previous methods for building attack graphs struggle with limited generalization and require expert model design. The AttacKG+ framework [324] addresses these challenges using four modules: rewriter, parser, identifier, and summarizer, powered by LLMs. The rewriter filters redundant information and organizes content into sections aligned

with MITRE TTP stages. The parser extracts a behavior graph containing atomic events, temporal relations, and entity links. The identifier maps this graph to MITRE technique labels. Finally, the summarizer provides concise summaries for each tactical stage, ensuring clarity. Despite its strengths, the AttacKG+ framework currently exhibits certain limitations. In particular, it struggles to effectively identify threat-related information from documents containing multiple modalities. Additionally, its ability to cross-reference and integrate information across various documents related to different Advanced Persistent Threats (APTs) is currently limited.

Language Models (LMs) have shown immense potential for automating the analysis of security logs, LogPrécis [325], leverages pre-trained language models to analyze Unix shell attack logs by mapping raw commands into intermediate representations that reflect attackers' objectives, guided by the MITRE ATT&CK framework. This approach produces high-level attack fingerprints, reducing nearly 400,000 unique shell logs to 3,000 distinct patterns, significantly simplifying forensic analysis, enhancing novelty detection, and identifying attack families. By employing few-shot fine-tuning on fewer samples, the tool effectively assigns tactics to log segments, enabling security analysts to understand attackers' goals and track variations over time. Attack techniques (AT) represent a higher-level intelligence that identifies consistent patterns in attack behaviors. AECR [326], built on a fine-tuned ChatGLM3-6B model, automates the extraction of AT intelligence by processing unstructured CTI reports, identifying AT-relevant sentences, and categorizing AT types. To bridge domain-specific knowledge gaps, ChatGLM is fine-tuned with P-Tuning v2, which enhances prefix-tuning by adding prompt tokens, increasing task-specific parameters, and improving performance. With its lightweight design and efficiency, AECR surpasses other LLMs like LLaMA and Qwen, demonstrating ChatGLM's superiority in extracting AT intelligence.

The study [327] evaluates the effectiveness of open-source and commercial LLM-based chatbots in recognizing cybersecurity-related tweets and extracting relevant information. Chatbots like ChatGPT, GPT4all, Dolly, Stanford Alpaca, Falcon, and Vicuna were tested for binary classification and NER tasks using OSINT data. The evaluation included normal datasets, shuffled datasets, and isolated prompt testing to ensure unbiased results. For NER tasks, methods like Entity-Specific Prompting (ESP) and Guide-Line Prompting (GLP) were tested, focusing on extracting specific entities such as organization names and product versions. Despite achieving competitive performance in binary classification, LLM-based chatbots still lag behind specialized models in accurately extracting structured cybersecurity information from tweets, highlighting areas for further improvement. Recently, SecLLMHolmes [328], an evaluation framework, was developed to assess the ability of LLMs to detect security vulnerabilities. Through 228 code scenarios across eight critical vulnerabilities in C and Python, the framework tested eight advanced LLMs. The results revealed significant limitations, including inconsistent outputs, high false positives, and flawed reasoning. Even leading models like PaLM2 and GPT-4 struggled with minor code changes, such as renaming variables or adding library functions, and failed to perform reliably in real-world scenarios. These findings highlight the need for further improvements, with SecLLMHolmes providing a valuable benchmark for guiding future advancements in LLM-based vulnerability detection.

LLMs like GPT-3 demonstrate both beneficial and malicious potential, particularly in generating text and code. Recent research [329] highlights how attackers can exploit GPT-3's capabilities to create functional malware through modular building blocks rather than complete descriptions. The study reports that GPT enables the generation of numerous malware variants, with detection rates varying from 4 to 55 on VirusTotal. While GPT-3 struggles with crafting long, complex malware routines in one step, it excels in breaking down tasks into smaller, manageable snippets. These snippets often include diverse API calls, allowing attackers to bypass API-based detection systems. Beyond malware generation, GPT-3's transformation abilities were used for

source code obfuscation, to evade detection. LLMs can be exploited for adversarial purposes, enabling attackers to craft highly convincing phishing emails that evade traditional detection systems. A recent study [330] evaluates the vulnerability of phishing detectors, such as Gmail Spam Filter and SpamAssassin, along with machine learning models, against these advanced threats. The study reveals a significant drop in detection accuracy when faced with LLM-rephrased phishing emails, emphasizing critical weaknesses in current defenses. Despite their misuse, LLMs also present opportunities for improving detection capabilities by generating diverse phishing datasets. For example, constructing datasets that combine original and rephrased phishing emails to train detectors more effectively. By exposing models to broader linguistic variations, these datasets enhance resilience against sophisticated attacks. The findings of this study demonstrate the dual role of LLMs in both advancing phishing attacks and transforming Cyber Threat Intelligence (CTI) through improved detection systems.

Lessons learned. The integration of LLMs with CTI has demonstrated significant advancements in automating knowledge extraction, threat detection, and attack attribution. Studies have leveraged LLMs such as GPT, Llama 2, Mistral 7B, and ChatGLM for structured knowledge graph construction, TTP extraction, and threat prediction. Techniques like few-shot learning, fine-tuning, and entity-specific prompting have improved the ability of LLMs to process unstructured threat intelligence. However, limitations persist in cross-referencing multimodal data, handling adversarial text modifications, and detecting vulnerabilities in real-world scenarios. Additionally, LLMs exhibit dual-use potential, as they not only improve CTI but also aid attackers in malware generation, phishing evasion, and API bypassing. The findings highlight the need for better prompt engineering, larger annotated datasets, and improved link prediction models to enhance CTI automation while mitigating adversarial exploitation of LLMs.

12. Challenges and open research directions

In this section, we delve into the primary challenges that researchers and cybersecurity professionals encounter in CTI, recognizing the complexities and intricacies involved in safeguarding our digital world. These challenges range from the handling of vast and diverse data sources to preserving privacy and from ensuring data quality to addressing the constant evolution of threat tactics.

12.1. Handling unstructured data

CTI data primarily originates from sources like social media, blogs, and threat reports. This data is often unstructured, making it essential to efficiently analyze, aggregate, and connect various data streams. To gather this data, Web crawlers are employed, and the researchers [32, 257] utilized multiple crawlers for this purpose. However, implementing Web crawlers for data collection poses challenges due to the technical complexities and the ever-changing nature of online content. Moreover, many threat reports contain extraneous and irrelevant information, making it challenging to extract the essential details related to attack behaviors. To address this issue, automatic solutions are needed. The researchers [59,167] in their study employed supervised classifiers to extract pertinent data, a process entailing the manual annotation of CTI data. Manual data annotation is both time-consuming and resource-intensive, and this becomes particularly challenging due to the dynamic nature of cyber threats, which demands continuous annotation efforts. Additionally, certain studies [44] incorporated transformer-based models for classification, including cybersecurity-specific models like SecBERT, CySecBERT, and CyBERT. Nevertheless, there exists a space for research to develop models capable of effectively handling emerging cybersecurity attack terminology.




Fig. 13. Literature timeline.

12.2. Adversarial data injection

As data primarily originates from open sources, it is imperative to acknowledge the potential threat posed by malicious actors who might deliberately inject poisoned or noisy data to disrupt the integrity of the CTI system. The impact of such fake data can be profound, affecting not only the quality of CTI but also the efficacy of threat mitigation strategies. However, only a limited number of studies have undertaken experiments to assess the impact of adversarial attacks [78] on CTI applications and the development of corresponding defense mechanisms [320]. This highlights a notable deficiency within the CTI field, emphasizing the urgent need for robust defense mechanisms capable of shielding against poisoned data introduced by adversaries during adversarial attacks.

12.3. Standardized quality assessment

In CTI, several quality challenges have emerged, highlighting critical research directions for the future. One key challenge is the overwhelming volume of data in the digital realm, making it difficult to distinguish genuine threats from noise. Additionally, the reliance on open-source information can introduce issues of verification and accuracy, potentially leading to false alarms or missed threats. The timely nature of threat intelligence is paramount, as threat actors constantly evolve, rendering outdated information obsolete and potentially leaving systems vulnerable. Addressing bias in threat intelligence sources is another critical research direction, requiring methods for source credibility assessment. To address these challenges, ongoing monitoring, data verification, and the development of standardized quality assessment methods are necessary, ultimately enhancing the overall value of CTI for informed decision-making and proactive cybersecurity defense.

12.4. Privacy and ethical considerations

When dealing with CTI, several important considerations come into play, particularly regarding privacy and legal issues related to data sharing. To safeguard cyber infrastructure against intentional threats, a collaborative effort involving cybersecurity professionals, researchers, and various organizations is essential. However, organizations may be hesitant to share information due to concerns about potential reputational damage resulting from disclosing details of cyberattacks. To

Table 12
Amount of papers analyzed per topic.

Topic	Amount of papers
Preprocessing and Tokenization	14
Text Representation Techniques	14
NLP Algorithms and Models	9
NLP Library and Tools	19
Crawling from Clear Web	16
Crawling from Social Media	26
Crawling from Dark/Deep Web	28
Text Classification	24
Text Similarity and Clustering	11
Text Mining and Open Information Extraction	5
Cross-Lingual Threat Intelligence	5
Topic Detection and Trend Analysis	11
Text Summarization	4
Named Entity Recognition	33
Event Identification	11
Relation Extraction	31
Knowledge Representation using Graphs	35
Explainability in CTI	4
CTI Information-sharing platforms	12
Standardization efforts	13
Quality and Privacy in CTI Sharing	20
Adversarial Attack	5
Role of Large Language Models in CTI	10

address these concerns, initiatives like TISP provide preliminary trust-building mechanisms, but these are currently limited to group-based access control and ranking systems [17]. Researchers are actively exploring techniques to enable organizations to share threat intelligence while preserving the confidentiality of sensitive information, employing approaches such as federated learning and Blockchain [300,302,304]. Federated learning, which facilitates collaborative ML model development without sharing raw data, holds promise for enhancing collective threat detection while upholding data privacy. Another noteworthy strategy, differential privacy, ensures that the inclusion or exclusion of specific data points has minimal impact on overall outcomes, making it a viable option for safeguarding individual privacy while extracting valuable insights from CTI. Furthermore, researchers are developing sophisticated access control mechanisms that restrict access to specific CTI data to authorized entities, employing fine-grained access controls and consent management. Nevertheless, further research is essential to foster trust in CTI sharing, exploring methods to establish confidence among sharing entities and ensuring the responsible and defensive use of shared data while mitigating the risk of misuse.

12.5. Multilingual NLP for threat intelligence

Enhancing the capabilities of NLP models to effectively analyze and understand threat intelligence across language barriers remains a pressing challenge. While English is commonly used, it is essential to recognize that the Dark Web, a hotspot for cyber threats, hosts content in a multitude of languages. Notably, there is a limited amount of research that delves into the linguistic analysis of non-English content on the Dark Web [143,162]. The challenges are exacerbated when considering text written in languages like Russian but using English characters. This linguistic diversity poses unique challenges in deciphering and extracting meaningful threat intelligence from non-English text. The multilingual nature of the Dark Web underscores the importance of cross-lingual threat analysis, which remains an active area for open research.

13. Summary and conclusion

Cyber Threats Intelligence is made up of a powerful set of tools offering valuable insights into possible cyber threats and attacks, empowering the adoption of proactive defense measures. The continuous

practice of CTI demands vigilant monitoring and constant adjustments to keep pace with the ever-changing threat environment. For this reason, recently, advanced ML methods have been applied to this field. In particular, NLP techniques have come in aid to generate robust and actionable intelligence when it comes to data collection from heterogeneous and textual data sources, data pre-processing, extraction of features, classification, clustering, and entity recognition. To give a contribution in this setting, in this survey, we provided a detailed review of the most significant works focusing on NLP-based techniques for CTI. By our proposed selection criteria, we thoroughly review and discuss the most relevant and recent works, including the ones dealing with the definitions and principles of CTI, the main data crawling techniques to extract data from main CTI Web sources, and NLP's models, techniques, and applications within the CTI domain. Furthermore, the survey also addresses the topics of CTI standardization protocols and results sharing. Finally, we discussed the main challenges and promising future research directions. In summary, we analyzed 250 articles published in renowned international conferences, journals, symposiums, and workshops with a focus on NLP-based techniques for CTI or related areas. Table 12 depicts a quantitative overview of the reviewed literature divided into topics, whereas Fig. 13 visualizes the analyzed number of articles published per year.

The study of the research works explored in this paper can be regarded as a foundation, as we plan to continue our investigation by diving into particular aspects only mentioned in this survey. For instance, an interesting direction can be the review of the paper exploiting existing CTI industrial platforms to give the reader a larger spectrum of diverse problems about them. Moreover, an extensive and exhaustive description of all the Authoritative Cybersecurity Databases available online is also a demanding task. Moreover, at the time of writing, LLMs have started to become a core NLP tool for diverse applications. However, to the best of our knowledge, proposals adopting LLMs to improve CTI management are still at a very early stage. We sincerely hope that this piece of work can aid both researchers and

practitioners in capturing the pivotal elements of this domain, clarifying the most significant progress, and highlighting future research advancements.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported in part by the following projects:

(i) The HORIZON Europe Framework Programme partly supported this work through the project "OPTIMA - Organization sPecific Threat Intelligence Mining and sharing, Italy" (101063107). Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or granting authority. Neither the European Union nor the granting authority can be held responsible for them.

(ii) The project SERICS, Italy (PE00000014) under the NRRP MUR program funded by the EU - NGEU. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the Italian MUR. Neither the European Union nor the Italian MUR can be held responsible for them.

(iii) The PRIN 2022 Project, Italy "HOMEY: a Human-centric IoE-based Framework for Supporting the Transition Towards Industry 5.0" (code: 2022NX7WKE, CUP: F53D23004340006) funded by the European

Union - Next Generation EU.



Funded by
the European Union



Finanziato
dall'Unione europea
NextGenerationEU

Appendix

See Table A.13

Table A.13

List of reference links for different libraries, tools, and threat intelligence platforms discussed in the paper.

Term	Source
Stanford CoreNLP	https://stanfordnlp.github.io/CoreNLP/
Python NLTK	https://www.nltk.org/
spaCy	https://spacy.io/
Flair	https://huggingface.co/flair/pos-english
Huggingface	https://huggingface.co/
ACHE crawler	https://github.com/ViDA-NYU/ache
SMILE	https://haifengl.github.io
JSoup Library	http://www.jsoup.org/
Rhino	https://developer.mozilla.org/docs/Mozilla/Projects/Rhino/
HTML Unit	http://htmlunit.sourceforge.net/
Beautiful Soup	https://www.crummy.com/software/BeautifulSoup/
Twitterscraper	https://pypi.org/project/twitterscraper/0.2.7/
APTNotes	https://github.com/kbandla/APTnotesarchived
MITRE ATT&CK	https://attack.mitre.org/
APT Groups and Operations	https://airtable.com/shr3Po3DsZUQZY4we/
XianZhi platform	http://service-corp.odps.aliyun-inc.com/api
CVE database	https://cve.mitre.org/
Common weakness enumeration (CWE)	https://cwe.mitre.org/
CAPEC database	https://capec.mitre.org/
WASC	http://www.webappsec.org/
NVD	https://nvd.nist.gov/
PacketStorm	https://packetstormsecurity.com/news/tags/database/
Privacy Rights Clearinghouse	https://www.privacyrights.org/
Hackmageddon	https://www.hackmageddon.com/
Tweepy	https://www.tweepy.org/
GetOldTweets	https://github.com/Jefferson-Henrique/GetOldTweets-java
GNIP's decahose	https://www.cabinetm.com/product/gnip/decahose
Reddit	https://www.reddit.com/

(continued on next page)

Table A.13 (continued).

Term	Source
GitHub	https://github.com/
Pastebin	https://pastebin.com/
PasteFS	https://pasteFS.com/
Pastelink	https://pastelink.net/
Wilders Security Forums	https://www.wilderssecurity.com/threads/
Oracle Security Blog	https://blogs.oracle.com/security/
Security Forum	https://www.securityforum.org/events/
Onion Router	https://www.torproject.org/
Chrome browser puppeteer	https://pptr.dev
Nulled.IO	https://www.nulled.to
Nulled.IO database	http://leakforums.net/thread-719337
vBulletin	https://www.vbulletin.com/
Scrapy	https://scrapy.org/
Sixgill	https://www.cybersixgill.com/
Wordnet	https://wordnet.princeton.edu/
OpenHowNet	https://github.com/thunlp/OpenHowNet
Russnet	http://www.russnet.org/
CoNLL-2003	https://paperswithcode.com/dataset/conll-2003
pyparsing	https://pypi.org/project/pyparsing/
TextRazor	https://www.textrazor.com/
BRON dataset	http://bron.alfa.csail.mit.edu/info.html
VirusTotal	https://www.virustotal.com/
ThreatFox	https://threatfox.abuse.ch/
URLHaus	https://urlhaus.abuse.ch/
MalwareBazaar	https://bazaar.abuse.ch/
YARAify	https://yaraify.abuse.ch/
ThreatStream	https://api.threatstream.com/
ThreatQ	https://www.threatq.com/
ThreatConnect	https://threatconnect.com/
Open Threat Exchange (OTX)	https://otx.alienvault.com/dashboard/new
IBM X-Force Exchange	https://exchange.xforce.ibmcloud.com/
Falcon X Intelligence CrowdStrike	https://www.cosive.com/capabilities/crowdstrike-falcon-x
Collective Intelligence Framework	https://www.mitre.org/our-impact/intellectual-property/collaborative-research-threats-crits
Anomali STAXX	https://www.anomali.com/resources/staxx
OpenCTI	https://demo.opencti.io/
STIX Web site	https://oasis-open.github.io/cti-documentation/examples/identifying-a-threat-actor-profile
TAXII	https://oasis-open.github.io/cti-documentation/taxii/intro
CyBOX	https://cybox.mitre.org/about/
MAEC	http://maecproject.github.io/about-maec/
VERIS	https://verisframework.org/

Data availability

No data was used for the research described in the article.

References

- [1] R. McMillan, Definition: Threat intelligence, 2022, <https://gartner.com/>.
- [2] Dave Shackelford, Who's Using Cyberthreat Intelligence and How, SANS Institute, 2015.
- [3] Henry Dalziel, How to Define and Build an Effective Cyber Threat Intelligence Capability, Syngress, 2014.
- [4] Giuseppe Cascavilla, Damian A. Tamburri, Willem-Jan Van Den Heuvel, Cybercrime threat intelligence: A systematic multi-vocal literature review, *Comput. Secur.* 105 (2021) 102258.
- [5] Md Rayhanur Rahman, Rezvan Mahdavi Hezaveh, Laurie Williams, What are the attackers doing now? Automating cyberthreat intelligence extraction from text on pace with the changing threat landscape: A survey, *ACM Comput. Surv.* 55 (12) (2023) 1–36.
- [6] Reza Montasari, Fiona Carroll, Stuart Macdonald, Hamid Jahankhani, Amin Hosseinian-Far, Alireza Daneshkhan, Application of artificial intelligence and machine learning in producing actionable cyber threat intelligence, *Digit. Forensic Investig. Internet Things (IoT) Devices* (2021) 47–64.
- [7] Amani Ibrahim, Dhananjay Thiruvady, Jean-Guy Schneider, Mohamed Abdelrazek, The challenges of leveraging threat intelligence to stop data breaches, *Front. Comput. Sci.* 2 (2020) 36.
- [8] Sagar Samtani, Murat Kantarcioglu, Hsinchun Chen, Trailblazing the artificial intelligence for cybersecurity discipline: a multi-disciplinary research roadmap, *ACM Trans. Manag. Inf. Syst. (TMIS)* 11 (4) (2020) 1–19.
- [9] Sagar Samtani, Maggie Abate, Victor Benjamin, Weifeng Li, Cybersecurity as an industry: A cyber threat intelligence perspective, in: *The Palgrave Handbook of International Cybercrime and Cyberdeviance*, Springer, 2020, pp. 135–154.
- [10] Randa Basheer, Bassel Alkhatib, Threats from the dark: A review over dark web investigation research for cyber threat intelligence, *J. Comput. Netw. Commun.* 2021 (2021) 1–21.
- [11] Tanja Miloshevska, Dark web as a contemporary challenge to cyber security, *Kriminalističke teme—Časopis za kriminalistiku, kriminologiju i sigurnosne studije* 19 (5) (2019) 117–128.
- [12] Chen Gao, Xuan Zhang, Mengting Han, Hui Liu, A review on cyber security named entity recognition, *Front. Inf. Technol. Electron. Eng.* 22 (9) (2021) 1153–1168.
- [13] Tiberiu-Marian Georgescu, Bogdan Iancu, Alin Zamfiroiu, Mihai Doinea, Catalin Emilian Boja, Cosmin Cartas, A survey on named entity recognition solutions applied for cybersecurity-related text processing, in: *Proceedings of Fifth International Congress on Information and Communication Technology*, Vol. 2, ICICT 2020, London, Springer, 2021, pp. 316–325.
- [14] Wiem Tounsi, Helmi Rais, A survey on technical threat intelligence in the age of sophisticated cyber attacks, *Comput. Secur.* 72 (2018) 212–233.
- [15] Ahmed El-Kosairy, Nashwa Abdelbaki, Heba Aslan, A survey on cyber threat intelligence sharing based on Blockchain, *Adv. Comput. Intell.* 3 (3) (2023) 10.
- [16] Thomas D. Wagner, Khaled Mahbub, Esther Palomar, Ali E. Abdallah, Cyber threat intelligence sharing: Survey and research directions, *Comput. Secur.* 87 (2019) 101589.
- [17] Md Sahrom Abu, Siti Rahayu Selamat, Aswami Ariffin, Robiah Yusof, Cyber threat intelligence—issue and challenges, *Indones. J. Electr. Eng. Comput. Sci.* 10 (1) (2018) 371–379.
- [18] Clemens Sauerwein, Daniel Fischer, Milena Rubsamén, Guido Rosenberger, Dirk Stelzer, Ruth Breu, From threat data to actionable intelligence: an exploratory analysis of the intelligence cycle implementation in cyber threat intelligence

- sharing platforms, in: Proceedings of the 16th International Conference on Availability, Reliability and Security, 2021, pp. 1–9.
- [19] Daniel Schlette, Marco Caselli, Günther Pernul, A comparative study on cyber threat intelligence: The security incident response perspective, *IEEE Commun. Surv. Tutor.* 23 (4) (2021) 2525–2556.
- [20] Nan Sun, Jun Zhang, Paul Rimba, Shang Gao, Leo Yu Zhang, Yang Xiang, Data-driven cybersecurity incident prediction: A survey, *IEEE Commun. Surv. Tutor.* 21 (2) (2018) 1744–1772.
- [21] Nan Sun, Ming Ding, Jiaojiao Jiang, Weikang Xu, Xiaoxing Mo, Yonghang Tai, Jun Zhang, Cyber threat intelligence mining for proactive cybersecurity defense: A survey and new perspectives, *IEEE Commun. Surv. Tutor.* (2023).
- [22] Poopak Alaeifar, Shantanu Pal, Zahra Jadidi, Mukhtar Hussain, Ernest Foo, Current approaches and future directions for cyber threat intelligence sharing: A survey, *J. Inf. Secur. Appl.* 83 (2024) 103786.
- [23] Dincy R. Arikkat, P. Vinod, K.A. Rafidha Rehiman, Serena Nicolazzo, Antonino Nocera, Mauro Conti, Relation extraction techniques in cyber threat intelligence, in: International Conference on Applications of Natural Language to Information Systems, Springer, 2024, pp. 348–363.
- [24] Yiren Chen, Mengjiao Cui, Ding Wang, Yiyang Cao, Peian Yang, Bo Jiang, Zhigang Lu, Baoxu Liu, A survey of large language models for cyber threat detection, *Comput. Secur.* (2024) 104016.
- [25] Google, Google scholar, 2023, <https://scholar.google.com>.
- [26] Elsevier, Scopus, 2023, <https://www.scopus.com>.
- [27] Scimago Lab, Scimago, 2023, <https://www.scimagojr.com/>.
- [28] Computing Research & Education, CORE Conference Portal, 2023, <https://portal.core.edu.au/conf-ranks>.
- [29] Ehtsham Irshad, Abdul Basit Siddiqui, Cyber threat attribution using unstructured reports in cyber threat intelligence, *Egypt. Inform. J.* 24 (1) (2023) 43–59.
- [30] Scott Ainslie, Dean Thompson, Sean Maynard, Atif Ahmad, Cyber-threat intelligence for security decision-making: A review and research agenda for practice, *Comput. Secur.* (2023) 103352.
- [31] Isuf Deliu, Carl Leichter, Katrin 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, IEEE, 2018, pp. 5008–5013.
- [32] Jun Zhao, Qiben Yan, Jianxin Li, Minglai Shao, Zuti He, Bo Li, TIMiner: Automatically extracting and analyzing categorized cyber threat intelligence from social data, *Comput. Secur.* 95 (2020) 101867.
- [33] Vahid Behzadan, Carlos Aguirre, Avishek Bose, William Hsu, Corpus and deep learning classifier for collection of cyber threat indicators in twitter stream, in: 2018 IEEE International Conference on Big Data, Big Data, IEEE, 2018, pp. 5002–5007.
- [34] Edward Loper, Steven Bird, Nltk: The natural language toolkit, 2002, arXiv preprint [Cs/0205028](https://arxiv.org/abs/0205028).
- [35] Vittorio Orbinato, Mariarosaria Barbaraci, Roberto Natella, Domenico Cotroneo, Automatic mapping of unstructured cyber threat intelligence: An experimental study: (practical experience report), in: 2022 IEEE 33rd International Symposium on Software Reliability Engineering, ISSRE, IEEE, 2022, pp. 181–192.
- [36] Peng Gao, Fei Shao, Xiaoyuan Liu, Xusheng Xiao, Zheng Qin, Fengyuan Xu, Prateek Mittal, Sanjeev R. Kulkarni, Dawn Song, Enabling efficient cyber threat hunting with cyber threat intelligence, in: 2021 IEEE 37th International Conference on Data Engineering, ICDE, IEEE, 2021, pp. 193–204.
- [37] Peng Gao, Fei Shao, Xiaoyuan Liu, Xusheng Xiao, Haoyuan Liu, Zheng Qin, Fengyuan Xu, Prateek Mittal, Sanjeev R. Kulkarni, Dawn Song, A system for efficiently hunting for cyber threats in computer systems using threat intelligence, in: 2021 IEEE 37th International Conference on Data Engineering, ICDE, IEEE, 2021, pp. 2705–2708.
- [38] Masashi Kadoguchi, Hanae Kobayashi, Shota Hayashi, Akira Otsuka, Masaki Hashimoto, Deep self-supervised clustering of the dark web for cyber threat intelligence, in: 2020 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2020, pp. 1–6.
- [39] Wenzhuo Yang, Kwok-Yan Lam, Automated cyber threat intelligence reports classification for early warning of cyber attacks in next generation SOC, in: Information and Communications Security: 21st International Conference, ICICS 2019, Beijing, China, December 15–17, 2019, Revised Selected Papers 21, Springer, 2020, pp. 145–164.
- [40] Taoran Ji, Xuchao Zhang, Nathan Self, Kaiqun Fu, Chang-Tien Lu, Naren Ramakrishnan, Feature driven learning framework for cybersecurity event detection, in: Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, 2019, pp. 196–203.
- [41] Hsin-Ju Chan, Chin-Yuan Hsu, Ching-Chang Chien, Ji-Jie Wu, He-Lin Ku, FeedRef2022: A named entity recognition dataset for extracting indicators of compromise, in: 2022 IEEE International Conference on Big Data, Big Data, IEEE, 2022, pp. 2578–2584.
- [42] Jacob Devlin, Ming-Wei Chang, Kenton Lee, Kristina Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, 2018, arXiv preprint [arXiv:1810.04805](https://arxiv.org/abs/1810.04805).
- [43] Pavlos Evangelatos, Christos Iliou, Thanassis Mavropoulos, Konstantinos Apostolou, Theodora Tsikrika, Stefanos Vrochidis, Ioannis Kompatsiaris, Named entity recognition in cyber threat intelligence using transformer-based models, in: 2021 IEEE International Conference on Cyber Security and Resilience, CSR, IEEE, 2021, pp. 348–353.
- [44] Philipp Kuehn, Mike Schmidt, Christian Reuter, ThreatCrawl: A BERT-based focused crawler for the cybersecurity domain, 2023, arXiv preprint [arXiv:2304.11960](https://arxiv.org/abs/2304.11960).
- [45] Soham Dasgupta, Aritran Piplai, Anantaa Kotal, Anupam Joshi, A comparative study of deep learning based named entity recognition algorithms for cybersecurity, in: 2020 IEEE International Conference on Big Data, Big Data, IEEE, 2020, pp. 2596–2604.
- [46] Xiaodi Wang, Jiayong Liu, A novel feature integration and entity boundary detection for named entity recognition in cybersecurity, *Knowl.-Based Syst.* 260 (2023) 110114.
- [47] Yin Zhang, Rong Jin, Zhi-Hua Zhou, Understanding bag-of-words model: a statistical framework, *Int. J. Mach. Learn. Cybern.* 1 (2010) 43–52.
- [48] Tomas Mikolov, Kai Chen, Gregory S. Corrado, Jeffrey Dean, Efficient estimation of word representations in vector space, in: International Conference on Learning Representations, 2013.
- [49] Juan Ramos, et al., Using tf-idf to determine word relevance in document queries, in: Proceedings of the First Instructional Conference on Machine Learning, vol. 242, no. 1, Citeseer, 2003, pp. 29–48.
- [50] Jeffrey Pennington, Richard Socher, Christopher D. Manning, GloVe: Global vectors for word representation, in: Empirical Methods in Natural Language Processing, EMNLP, 2014, pp. 1532–1543.
- [51] Sampashree Nayak, Deepak Nadig, Byrav Ramamurthy, Analyzing malicious URLs using a threat intelligence system, in: 2019 IEEE International Conference on Advanced Networks and Telecommunications Systems, ANTS, IEEE, 2019, pp. 1–4.
- [52] Benjamin Ampel, Sagar Samtani, Hongyi Zhu, Steven Ullman, Hsinchun Chen, Labeling hacker exploits for proactive cyber threat intelligence: a deep transfer learning approach, in: 2020 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2020, pp. 1–6.
- [53] Masashi Kadoguchi, Shota Hayashi, Masaki Hashimoto, Akira Otsuka, Exploring the dark web for cyber threat intelligence using machine learning, in: 2019 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2019, pp. 200–202.
- [54] Nuno Dionísio, Fernando Alves, Pedro M. Ferreira, Alysso Bessani, Cyberthreat detection from twitter using deep neural networks, in: 2019 International Joint Conference on Neural Networks, IJCNN, IEEE, 2019, pp. 1–8.
- [55] Gyeongmin Kim, Chanhee Lee, Jaechoon Jo, Heuseok Lim, Automatic extraction of named entities of cyber threats using a deep Bi-LSTM-CRF network, *Int. J. Mach. Learn. Cybern.* 11 (2020) 2341–2355.
- [56] Ba Dung Le, Guanhua Wang, Mehwish Nasim, Ali Babar, Gathering cyber threat intelligence from Twitter using novelty classification, 2019, arXiv preprint [arXiv:1907.01755](https://arxiv.org/abs/1907.01755).
- [57] Victor Adewopo, Bilal Gonen, Festus Adewopo, Exploring open source information for cyber threat intelligence, in: 2020 IEEE International Conference on Big Data, Big Data, IEEE, 2020, pp. 2232–2241.
- [58] Tao Bo, Yue Chen, Can Wang, Yunwei Zhao, Kwok-Yan Lam, Chi-Hung Chi, Hui Tian, Tom: A threat operating model for early warning of cyber security threats, in: Advanced Data Mining and Applications: 15th International Conference, ADMA 2019, Dalian, China, November 21–23, 2019, Proceedings 15, Springer, 2019, pp. 696–711.
- [59] Linn-Mari Kristiansen, Vinti Agarwal, Katrin Franke, Raj Sanjay Shah, CTI-Twitter: gathering cyber threat intelligence from twitter using integrated supervised and unsupervised learning, in: 2020 IEEE International Conference on Big Data, Big Data, IEEE, 2020, pp. 2299–2308.
- [60] George Pantelis, Petros Petrou, Sophia Karagiorgou, Dimitrios Alexandrou, On strengthening smes and mes threat intelligence and awareness by identifying data breaches, stolen credentials and illegal activities on the dark web, in: Proceedings of the 16th International Conference on Availability, Reliability and Security, 2021, pp. 1–7.
- [61] Dong Li, Xiao Zhou, Ao Xue, Open source threat intelligence discovery based on topic detection, in: 2020 29th International Conference on Computer Communications and Networks, ICCCN, IEEE, 2020, pp. 1–4.
- [62] Gbadebo Ayoade, Swarup Chandra, Latifur Khan, Kevin Hamlen, Bhavani Thuraisingham, Automated threat report classification over multi-source data, in: 2018 IEEE 4th International Conference on Collaboration and Internet Computing, CIC, IEEE, 2018, pp. 236–245.
- [63] Katerina Goseva-Popstojanova, Jacob Tyo, Identification of security related bug reports via text mining using supervised and unsupervised classification, in: 2018 IEEE International Conference on Software Quality, Reliability and Security, QRS, IEEE, 2018, pp. 344–355.
- [64] Mengming Li, Rongfeng Zheng, Liang Liu, Pin Yang, Extraction of threat actions from threat-related articles using multi-label machine learning classification method, in: 2019 2nd International Conference on Safety Produce Informatization, IICSPI, IEEE, 2019, pp. 428–431.

- [65] Andrei Lima Queiroz, Susan McKeever, Brian Keegan, Eavesdropping hackers: Detecting software vulnerability communication on social media using text mining, in: *The Fourth International Conference on Cyber-Technologies and Cyber-Systems*, 2019, pp. 41–48.
- [66] Baidyanath Biswas, Arunabha Mukhopadhyay, Sudip Bhattacharjee, Ajay Kumar, Dursun Delen, A text-mining based cyber-risk assessment and mitigation framework for critical analysis of online hacker forums, *Decis. Support Syst.* 152 (2022) 113651.
- [67] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, Luke Zettlemoyer, Deep contextualized word representations, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, 2018.
- [68] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al., Improving language understanding by generative pre-training, 2018, OpenAI.
- [69] Leo Gao, John Schulman, Jacob Hilton, Scaling laws for reward model over-optimization, in: *International Conference on Machine Learning*, PMLR, 2023, pp. 10835–10866.
- [70] Tyna Eloundou, Sam Manning, Pamela Mishkin, Daniel Rock, Gpts are gpts: An early look at the labor market impact potential of large language models, 2023, arXiv preprint arXiv:2303.10130.
- [71] R. OpenAI, GPT-4 technical report, 2023, pp. 2303–08774, arXiv.
- [72] Yinghai Zhou, Yitong Ren, Ming Yi, Yanjun Xiao, Zhiyuan Tan, Nour Moustafa, Zhihong Tian, CDTier: A Chinese dataset of threat intelligence entity relationships, *IEEE Trans. Sustain. Comput.* (2023).
- [73] Markus Bayer, Philipp Kuehn, Ramin Shanehsaz, Christian Reuter, Cysecbert: A domain-adapted language model for the cybersecurity domain, *ACM Trans. Priv. Secur.* 27 (2) (2024) 1–20.
- [74] Mohamed Amine Ferrag, Mthandazo Ndhlovu, Norbert Tihanyi, Lucas C. Cordeiro, Merouane Debbah, Thierry Lestable, Narinderjit Singh Thandi, Revolutionizing cyber threat detection with large language models: A privacy-preserving bert-based lightweight model for IoT/IIoT devices, *IEEE Access* (2024).
- [75] Md Tanvirul Alam, Dipkamal Bhusal, Youngja Park, Nidhi Rastogi, Cyner: A Python library for cybersecurity named entity recognition, 2022, arXiv preprint arXiv:2204.05754.
- [76] Jiao Yin, MingJian Tang, Jinli Cao, Hua Wang, Apply transfer learning to cybersecurity: Predicting exploitability of vulnerabilities by description, *Knowl.-Based Syst.* 210 (2020) 106529.
- [77] Jian Liu, Junjie Yan, Jun Jiang, Yitong He, Xuren Wang, Zhengwei Jiang, Peian Yang, Ning Li, TriCTI: an actionable cyber threat intelligence discovery system via trigger-enhanced neural network, *Cybersecurity* 5 (1) (2022) 8.
- [78] Priyanka Ranade, Aritrnan Piplai, Sudip Mittal, Anupam Joshi, Tim Finin, Generating fake cyber threat intelligence using transformer-based models, in: *2021 International Joint Conference on Neural Networks, IJCNN, IEEE, 2021*, pp. 1–9.
- [79] Febrían Setianto, Erion Tsani, Fatima Sadiq, Georgios Domalis, Dimitris Tsakalidis, Panos Kostakos, GPT-2C: A parser for honeypot logs using large pre-trained language models, in: *Proceedings of the 2021 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2021, pp. 649–653.
- [80] Priyanka Ranade, Aritrnan Piplai, Anupam Joshi, Tim Finin, Cybert: Contextualized embeddings for the cybersecurity domain, in: *2021 IEEE International Conference on Big Data, Big Data, IEEE, 2021*, pp. 3334–3342.
- [81] Avishek Bose, Vahid Behzadan, Carlos Aguirre, William H. Hsu, A novel approach for detection and ranking of trendy and emerging cyber threat events in twitter streams, in: *Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, 2019, pp. 871–878.
- [82] Ying Dong, Wenbo Guo, Yueqi Chen, Xinyu Xing, Yuqing Zhang, Gang Wang, Towards the detection of inconsistencies in public security vulnerability reports, in: *28th USENIX Security Symposium, USENIX Security 19*, 2019, pp. 869–885.
- [83] Kiavash Satvat, Rigel Gjomemo, V.N. Venkatakrishnan, Extractor: Extracting attack behavior from threat reports, in: *2021 IEEE European Symposium on Security and Privacy, EuroS&P, IEEE, 2021*, pp. 598–615.
- [84] Nan Sun, Jun Zhang, Shang Gao, Leo Yu Zhang, Seyit Camtepe, Yang Xiang, Cyber information retrieval through pragmatics understanding and visualization, *IEEE Trans. Dependable Secur. Comput.* 20 (2) (2022) 1186–1199.
- [85] Paris Koloveas, Thanasis Chantzios, Christos Tryfonopoulos, Spiros Skiadopoulos, A crawler architecture for harvesting the clear, social, and dark web for IoT-related cyber-threat intelligence, in: *2019 IEEE World Congress on Services, Vol. 2642, SERVICES, IEEE, 2019*, pp. 3–8.
- [86] Paris Koloveas, Thanasis Chantzios, Sofia Alevizopoulou, Spiros Skiadopoulos, Christos Tryfonopoulos, intime: A machine learning-based framework for gathering and leveraging web data to cyber-threat intelligence, *Electronics* 10 (7) (2021) 818.
- [87] A.K. Singh, Navneet Goyal, Malcrawler: A crawler for seeking and crawling malicious websites, in: *Distributed Computing and Internet Technology: 13th International Conference, ICDCIT 2017, Bhubaneswar, India, January 13-16, 2017, Proceedings 13*, Springer, 2017, pp. 210–223.
- [88] Adrián Barbaresi, Trafilatura: A web scraping library and command-line tool for text discovery and extraction, in: *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, 2021, pp. 122–131.
- [89] Ke Li, Hui Wen, Hong Li, Hongsong Zhu, Limin Sun, Security OSIF: Toward automatic discovery and analysis of event based cyber threat intelligence, in: *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation, SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI, IEEE, 2018*, pp. 741–747.
- [90] Taneeya Satyapanich, Francis Ferraro, Tim Finin, Casie: Extracting cybersecurity event information from text, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8749–8757.
- [91] Umara Noor, Zahid Anwar, Tehmina Amjad, Kim-Kwang Raymond Choo, A machine learning-based FinTech cyber threat attribution framework using high-level indicators of compromise, *Future Gener. Comput. Syst.* 96 (2019) 227–242.
- [92] Chia-Mei Chen, Yu-Xuan Wang, Zheng-Xun Cai, Boyi Lee, Gu-Hsin Lai, Automatic summarization of critical threat intelligence using transfer learning, in: *International Computer Symposium*, Springer, 2022, pp. 343–348.
- [93] Chia-Mei Chen, Jing-Yun Kan, Ya-Hui Ou, Zheng-Xun Cai, Albert Guan, et al., Threat action extraction using information retrieval, in: *CS & IT Conference Proceedings*, vol. 11, no. 7, CS & IT Conference Proceedings, 2021.
- [94] Lior Perry, Bracha Shapira, Rami Puzis, No-doubt: Attack attribution based on threat intelligence reports, in: *2019 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2019*, pp. 80–85.
- [95] Yiming Wu, Qianjun Liu, Xiaojing Liao, Shouling Ji, Peng Wang, Xiaofeng Wang, Chunming Wu, Zhao Li, Price tag: towards semi-automatically discovery tactics, techniques and procedures of E-commerce cyber threat intelligence, *IEEE Trans. Dependable Secur. Comput.* (2021).
- [96] Dincy R. Arikkat, P. Vinod, K.A. Rafidha Rehiman, Serena Nicolazzo, Antonino Nocera, Georgiana Timpau, Mauro Conti, OSTIS: A novel organization-specific threat intelligence system, *Comput. Secur.* 145 (2024) 103990.
- [97] Nan Sun, Jun Zhang, Shang Gao, Leo Yu Zhang, Seyit Camtepe, Yang Xiang, Data analytics of crowdsourced resources for cybersecurity intelligence, in: *Network and System Security: 14th International Conference, NSS 2020, Melbourne, VIC, Australia, November 25–27, 2020, Proceedings 14*, Springer, 2020, pp. 3–21.
- [98] Fernando Alves, Ambrose Andongabo, Ilir Gashi, Pedro M. Ferreira, Alysso Bessani, Follow the blue bird: A study on threat data published on Twitter, in: *Computer Security—ESORICS 2020: 25th European Symposium on Research in Computer Security, ESORICS 2020, Guildford, UK, September 14–18, 2020, Proceedings, Part I 25*, Springer, 2020, pp. 217–236.
- [99] P.R. Vishnu, P. Vinod, Suleiman Y. Yerima, A deep learning approach for classifying vulnerability descriptions using self attention based neural network, *J. Netw. Syst. Manage.* 30 (2022) 1–27.
- [100] Trend Micro Incorporated, Hunting threats on Twitter: How social media can be used to gather actionable threat intelligence, 2023, <https://www.trendmicro.com/vinfo/es/security/news/cybercrime-and-digital-threats/hunting-threats-on-twitter>.
- [101] Dincy R. Arikkat, Andrea Di Sorbo, Corrado A. Visaggio, Mauro Conti, et al., Can Twitter be used to acquire reliable alerts against novel cyber attacks? 2023, arXiv preprint arXiv:2306.16087.
- [102] Fernando Alves, Aurélien Bettini, Pedro M. Ferreira, Alysso Bessani, Processing tweets for cybersecurity threat awareness, *Inf. Syst.* 95 (2021) 101586.
- [103] Ariel Rodriguez, Koji Okamura, Generating real time cyber situational awareness information through social media data mining, in: *2019 IEEE 43rd Annual Computer Software and Applications Conference, Vol. 2, COMPSAC, IEEE, 2019*, pp. 502–507.
- [104] Anna Sapienza, Sindhu Kiranmai Ernala, Alessandro Bessi, Kristina Lerman, Emilio Ferrara, Discover: Mining online chatter for emerging cyber threats, in: *Companion Proceedings of the the Web Conference 2018*, 2018, pp. 983–990.
- [105] Prasasthy Balasubramanian, Sadaf Nazari, Daniah Khosh Kholgh, Alireza Mahmoodi, Justin Seby, Panos Kostakos, A cognitive platform for collecting cyber threat intelligence and real-time detection using cloud computing, *Decis. Anal. J.* (2025) 100545.
- [106] U.S. Department of Homeland Security, Analyst’s desktop binder, 2011, <https://archive.epic.org/foia/epic-v-dhs-media-monitoring/Analyst-Desktop-Binder-REDACTED.pdf>.
- [107] Amirreza Niakanlahiji, Lida Safarnejad, Reginald Harper, Bei-Tsung Chu, Iocminer: Automatic extraction of indicators of compromise from twitter, in: *2019 IEEE International Conference on Big Data, Big Data, IEEE, 2019*, pp. 4747–4754.
- [108] Hyejin Shin, WooChul Shim, Saebom Kim, Sol Lee, Yong Goo Kang, Yong Ho Hwang, # twiti: Social listening for threat intelligence, in: *Proceedings of the Web Conference 2021*, 2021, pp. 92–104.
- [109] Xiuwen Liu, Jianming Fu, Yanjiao Chen, Event evolution model for cybersecurity event mining in tweet streams, *Inform. Sci.* 524 (2020) 254–276.

- [110] Ying Chen, J.D. Elenee Argentinis, Griff Weber, IBM Watson: how cognitive computing can be applied to big data challenges in life sciences research, *Clin. Ther.* 38 (4) (2016) 688–701.
- [111] Clemens Sauerwein, Christian Sillaber, Michael M. Huber, Andrea Mussmann, Ruth Brey, The tweet advantage: An empirical analysis of 0-day vulnerability information shared on twitter, in: *ICT Systems Security and Privacy Protection: 33rd IFIP TC 11 International Conference, SEC 2018, Held At the 24th IFIP World Computer Congress, WCC 2018, Poznan, Poland, September 18-20, 2018, Proceedings 33*, Springer, 2018, pp. 201–215.
- [112] Romilla Syed, Maryam Rahafruz, Jeffrey M. Keisler, What it takes to get retweeted: An analysis of software vulnerability messages, *Comput. Hum. Behav.* 80 (2018) 207–215.
- [113] Shin-Ying Huang, Tao Ban, Monitoring social media for vulnerability-threat prediction and topic analysis, in: *2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications, TrustCom, IEEE, 2020*, pp. 1771–1776.
- [114] Ambrose Andongabo, Ilir Gashi, vepRisk-A web based analysis tool for public security data, in: *2017 13th European Dependable Computing Conference, EDCC, IEEE, 2017*, pp. 135–138.
- [115] Sameera Horawalavithana, Abhishek Bhattacharjee, Renhao Liu, Nazim Choudhury, Lawrence O. Hall, Adriana Iamnitich, Mentions of security vulnerabilities on reddit, twitter and github, in: *IEEE/WIC/ACM International Conference on Web Intelligence, 2019*, pp. 200–207.
- [116] Yi-Ting Huang, R. Vaitheeshwari, Meng-Chang Chen, Ying-Dar Lin, Ren-Hung Hwang, Po-Ching Lin, Yuan-Cheng Lai, Eric Hsiao-Kuang Wu, Chung-Hsuan Chen, Zi-Jie Liao, et al., MITREtrieval: Retrieving MITRE techniques from unstructured threat reports by fusion of deep learning and ontology, *IEEE Trans. Netw. Serv. Manag.* (2024).
- [117] Valentine Legoy, Marco Caselli, Christin Seifert, Andreas Peter, Automated retrieval of ATT&CK tactics and techniques for cyber threat reports, 2020.
- [118] Tala Vahedi, Benjamin Ampel, Sagar Samtani, Hsinchun Chen, Identifying and categorizing malicious content on paste sites: a neural topic modeling approach, in: *2021 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2021*, pp. 1–6.
- [119] Rupinder Paul Khandpur, Taoran Ji, Steve Jan, Gang Wang, Chang-Tien Lu, Naren Ramakrishnan, Crowdsourcing cybersecurity: Cyber attack detection using social media, in: *Proceedings of the 2017 ACM on Conference on Information and Knowledge Management, 2017*, pp. 1049–1057.
- [120] Aldo Hernandez-Suarez, Gabriel Sanchez-Perez, Karina Toscano-Medina, Victor Martinez-Hernandez, Hector Perez-Meana, Jesus Olivares-Mercado, Victor Sanchez, Social sentiment sensor in twitter for predicting cyber-attacks using l1 regularization, *Sensors* 18 (5) (2018) 1380.
- [121] Ariel Rodriguez, Koji Okamura, Social media data mining for proactive cyber defense, *J. Inf. Process.* 28 (2020) 230–238.
- [122] Po-Yi Du, Ning Zhang, Mohammedreza Ebrahimi, Sagar Samtani, Ben Lazarine, Nolan Arnold, Rachael Dunn, Sandeep Suntwal, Guadalupe Angeles, Robert Schweitzer, et al., Identifying, collecting, and presenting hacker community data: Forums, IRC, carding shops, and DNMs, in: *2018 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2018*, pp. 70–75.
- [123] Victor Benjamin, Joseph S. Valacich, Hsinchun Chen, DICE-E: A framework for conducting darknet identification, collection, evaluation with ethics, *MIS Q.* 43 (1) (2019).
- [124] Mohammadreza Ebrahimi, Jay F. Nunamaker Jr., Hsinchun Chen, Semi-supervised cyber threat identification in dark net markets: A transductive and deep learning approach, *J. Manage. Inf. Syst.* 37 (3) (2020) 694–722.
- [125] Mohammed Almukaynizi, Ericsson Marin, Eric Nunes, Paulo Shakarian, Gerardo I. Simari, Dipsy Kapoor, Timothy Siedlecki, Darkment: A deployed system to predict enterprise-targeted external cyberattacks, in: *2018 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2018*, pp. 31–36.
- [126] Ryan Williams, Sagar Samtani, Mark Patton, Hsinchun 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, IEEE, 2018*, pp. 94–99.
- [127] Nazgol Tavabi, Nathan Bartley, Andres Abeliuk, Sandeep Soni, Emilio Ferrara, Kristina Lerman, Characterizing activity on the deep and dark web, in: *Companion Proceedings of the 2019 World Wide Web Conference, 2019*, pp. 206–213.
- [128] Sagar Samtani, Hongyi Zhu, Hsinchun Chen, Proactively identifying emerging hacker threats from the dark web: A diachronic graph embedding framework (d-gef), *ACM Trans. Priv. Secur. (TOPS)* 23 (4) (2020) 1–33.
- [129] Matthias Schäfer, Markus Fuchs, Martin Strohmeier, Markus Engel, Marc Liechti, Vincent Lenders, BlackWidow: Monitoring the dark web for cyber security information, in: *2019 11th International Conference on Cyber Conflict, Vol. 900, CyCon, IEEE, 2019*, pp. 1–21.
- [130] Nolan Arnold, Mohammadreza Ebrahimi, Ning Zhang, Ben Lazarine, Mark Patton, Hsinchun Chen, Sagar Samtani, Dark-net ecosystem cyber-threat intelligence (CTI) tool, in: *2019 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2019*, pp. 92–97.
- [131] Sagar Samtani, Yidong Chai, Hsinchun Chen, Linking exploits from the dark web to known vulnerabilities for proactive cyber threat intelligence: An attention-based deep structured semantic model, *MIS Q.* 46 (2) (2022).
- [132] Cheng Huang, Yongyan Guo, Wenbo Guo, Ying Li, HackerRank: Identifying key hackers in underground forums, *Int. J. Distrib. Sens. Netw.* 17 (5) (2021) 15501477211015145.
- [133] Hatma Suryotrisongko, Hari Ginardi, Henning Titi Ciptaningtyas, Saeed Dehqan, Yasuo Musashi, Topic modeling for Cyber Threat Intelligence (CTI), in: *2022 Seventh International Conference on Informatics and Computing, ICIC, IEEE, 2022*, pp. 1–7.
- [134] Tommaso Paladini, Lara Ferro, Mario Polino, Stefano Zanero, Michele Carminati, You might have known it earlier: Analyzing the role of underground forums in threat intelligence, in: *Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses, 2024*, pp. 368–383.
- [135] Sergio Pastrana, Daniel R. Thomas, Alice Hutchings, Richard Clayton, Crimebb: Enabling cybercrime research on underground forums at scale, in: *Proceedings of the 2018 World Wide Web Conference, 2018*, pp. 1845–1854.
- [136] Sagar Samtani, Kory Chinn, Cathy Larson, Hsinchun Chen, Azsecure hacker assets portal: Cyber threat intelligence and malware analysis, in: *2016 IEEE Conference on Intelligence and Security Informatics, ISI, IEEE, 2016*, pp. 19–24.
- [137] Isuf Deliu, Carl Leichter, Katrin 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, IEEE, 2017*, pp. 3648–3656.
- [138] Eric Nunes, Ahmad Diab, Andrew Gunn, Ericsson Marin, Vineet Mishra, Vivin Paliath, John Robertson, Jana Shakarian, Amanda Thart, Paulo Shakarian, Darknet and deepnet mining for proactive cybersecurity threat intelligence, in: *2016 IEEE Conference on Intelligence and Security Informatics, ISI, IEEE, 2016*, pp. 7–12.
- [139] Nazgol Tavabi, Palash Goyal, Mohammed Almukaynizi, Paulo Shakarian, Kristina Lerman, Darkembed: Exploit prediction with neural language models, in: *Proceedings of the AAAI Conference on Artificial Intelligence, vol. 32, no. 1, 2018*.
- [140] Ericsson Marin, Mohammed Almukaynizi, Eric Nunes, Jana Shakarian, Paulo Shakarian, Predicting hacker adoption on darkweb forums using sequential rule mining, in: *2018 IEEE Intl Conf on Parallel & Distributed Processing with Applications, Ubiquitous Computing & Communications, Big Data & Cloud Computing, Social Computing & Networking, Sustainable Computing & Communications, ISPA/IUCC/BDCloud/SocialCom/SustainCom, IEEE, 2018*, pp. 1183–1190.
- [141] Ericsson Marin, Jana Shakarian, Paulo Shakarian, Mining key-hackers on darkweb forums, in: *2018 1st International Conference on Data Intelligence and Security, ICDIS, IEEE, 2018*, pp. 73–80.
- [142] Eric Nunes, Paulo Shakarian, Gerardo I. Simari, At-risk system identification via analysis of discussions on the darkweb, in: *2018 APWG Symposium on Electronic Crime Research, ECrime, IEEE, 2018*, pp. 1–12.
- [143] Mohammadreza Ebrahimi, Mihai Surdeanu, Sagar Samtani, Hsinchun 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, IEEE, 2018*, pp. 85–90.
- [144] Fangzhou Dong, Shaoxian Yuan, Haoran Ou, Liang Liu, New cyber threat discovery from darknet marketplaces, in: *2018 IEEE Conference on Big Data and Analytics, ICBDA, IEEE, 2018*, pp. 62–67.
- [145] Sicong Shao, Cihan Tunc, Amany Al-Shawi, Salim Hari, Autonomic author identification in internet relay chat (IRC), in: *2018 IEEE/ACS 15th International Conference on Computer Systems and Applications, AICCSA, IEEE, 2018*, pp. 1–8.
- [146] Fernando Alves, Pedro Miguel Ferreira, Alysson Bessani, Design of a classification model for a twitter-based streaming threat monitor, in: *2019 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Workshops, DSN-W, IEEE, 2019*, pp. 9–14.
- [147] Nurfadilah Ariffini, Anazida Zainal, Mohd Aizaini Maarof, Mohamad Nizam Kassim, Ransomware entities classification with supervised learning for informal text, in: *2019 International Conference on Cybersecurity, ICocSec, IEEE, 2019*, pp. 86–90.
- [148] Paulo M.M.R. Alves, P.R. Geraldo Filho, Vinicius P. Gonçalves, Leveraging BERT's power to classify TTP from unstructured text, in: *2022 Workshop on Communication Networks and Power Systems, WCNPS, IEEE, 2022*, pp. 1–7.
- [149] Lingzi Li, Cheng Huang, Junren Chen, Automated discovery and mapping ATT&CK tactics and techniques for unstructured cyber threat intelligence, *Comput. Secur.* 140 (2024) 103815.
- [150] Ziyun Zhu, Tudor Dumitras, Chainsmith: Automatically learning the semantics of malicious campaigns by mining threat intelligence reports, in: *2018 IEEE European Symposium on Security and Privacy, EuroS&P, IEEE, 2018*, pp. 458–472.
- [151] Mariella Mischinger, Sergio Pastrana, Guillermo Suarez-Tangil, et al., IoC stalker: Early detection of indicators of compromise, in: *Annual Computer Security Applications Conference, 2024*.

- [152] Nanda Rani, Bikash Saha, Vikas Maurya, Sandeep Kumar Shukla, TTPXHunter: Actionable threat intelligence extraction as TTPs from finished cyber threat reports, *Digit. Threat.: Res. Pr.* 5 (4) (2024) 1–19.
- [153] Rui Azevedo, Ibéria Medeiros, Alysso Bessani, PURE: Generating quality threat intelligence by clustering and correlating OSINT, in: 2019 18th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/13th IEEE International Conference on Big Data Science and Engineering, *TrustCom/BigDataSE*, IEEE, 2019, pp. 483–490.
- [154] Yuma Kurogome, Yuto Otsuki, Yuhei Kawakoya, Makoto Iwamura, Syogo Hayashi, Tatsuya Mori, Koushik Sen, EIGER: automated IOC generation for accurate and interpretable endpoint malware detection, in: Proceedings of the 35th Annual Computer Security Applications Conference, 2019, pp. 687–701.
- [155] Sadegh M. Milajerdi, Birhanu Eshete, Rigel Gjomemo, V.N. Venkatakrishnan, Poirot: Aligning attack behavior with kernel audit records for cyber threat hunting, in: Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security, 2019, pp. 1795–1812.
- [156] Injy Sarhan, Marco Spruit, Open-cykg: An open cyber threat intelligence knowledge graph, *Knowl.-Based Syst.* 233 (2021) 107524.
- [157] Max Landauer, Florian Skopik, Markus Wurzenberger, Wolfgang Hotwagner, Andreas Rauber, A framework for cyber threat intelligence extraction from raw log data, in: 2019 IEEE International Conference on Big Data, *Big Data*, IEEE, 2019, pp. 3200–3209.
- [158] Tieming Chen, Huan Zeng, Mingqi Lv, Tiantian Zhu, CTIMD: cyber threat intelligence enhanced malware detection using API call sequences with parameters, *Comput. Secur.* 136 (2024) 103518.
- [159] Sagar Samtani, Ryan Chinn, Hsinchun Chen, Jay F. Nunamaker Jr., Exploring emerging hacker assets and key hackers for proactive cyber threat intelligence, *J. Manage. Inf. Syst.* 34 (4) (2017) 1023–1053.
- [160] Javid Ebrahimi, Daniel Lowd, Dejing Dou, On adversarial examples for character-level neural machine translation, in: Proceedings of the 27th International Conference on Computational Linguistics, 2018, pp. 653–663.
- [161] Mohammadreza Ebrahimi, Sagar Samtani, Yidong Chai, Hsinchun Chen, Detecting cyber threats in non-English hacker forums: an adversarial cross-lingual knowledge transfer approach, in: 2020 IEEE Security and Privacy Workshops, *SPW*, IEEE, 2020, pp. 20–26.
- [162] Priyanka Ranade, Sudip Mittal, Anupam Joshi, Karuna Joshi, Using deep neural networks to translate multi-lingual threat intelligence, in: 2018 IEEE International Conference on Intelligence and Security Informatics, *ISI*, IEEE, 2018, pp. 238–243.
- [163] Wanxiang Che, Zhenghua Li, Ting Liu, Ltp: A chinese language technology platform, in: *Coling 2010: Demonstrations*, 2010, pp. 13–16.
- [164] David M. Blei, Andrew Y. Ng, Michael I. Jordan, Latent dirichlet allocation, *J. Mach. Learn. Res.* 3 (Jan) (2003) 993–1022.
- [165] Tatsuya Nagai, Makoto Takita, Keisuke Furumoto, Yoshiaki Shiraishi, Kelin Xia, Yasuhiro Takano, Masami Mohri, Masakatu Morii, Understanding attack trends from security blog posts using guided-topic model, *J. Inf. Process.* 27 (2019) 802–809.
- [166] Jagadeesh Jagarlamudi, Hal Daumé III, Raghavendra Udapa, Incorporating lexical priors into topic models, in: Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics, 2012, pp. 204–213.
- [167] Md Imran Hossen, Ashrafal Islam, Farzana Anowar, Eshtiaq Ahmed, Mohammed Masudur Rahman, Generating cyber threat intelligence to discover potential security threats using classification and topic modeling, in: *Cyber Security using Modern Technologies*, CRC Press, 2021, pp. 141–153.
- [168] Maarten Grootendorst, BERTopic: Neural topic modeling with a class-based TF-IDF procedure, 2022, arXiv preprint arXiv:2203.05794.
- [169] Dima Angelov, Top2vec: Distributed representations of topics, 2020, arXiv preprint arXiv:2008.09470.
- [170] Jennifer Sleeman, Tim Finin, Milton Halem, Understanding cybersecurity threat trends through dynamic topic modeling, *Front. Big Data* 4 (2021) 601529.
- [171] Luheng He, Kenton Lee, Mike Lewis, Luke Zettlemoyer, Deep semantic role labeling: What works and what's next, in: Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), 2017, pp. 473–483.
- [172] Ernesto Rosario Russo, Andrea Di Sorbo, Corrado A. Visaggio, Gerardo Canfora, Summarizing vulnerabilities' descriptions to support experts during vulnerability assessment activities, *J. Syst. Softw.* 156 (2019) 84–99.
- [173] Casey Hanks, Michael Maiden, Priyanka Ranade, Tim Finin, Anupam Joshi, et al., Recognizing and extracting cybersecurity entities from text, in: Workshop on Machine Learning for Cybersecurity, International Conference on Machine Learning, 2022.
- [174] Aritr Piplai, Sudip Mittal, Anupam Joshi, Tim Finin, James Holt, Richard Zak, Creating cybersecurity knowledge graphs from malware after action reports, *IEEE Access* 8 (2020) 211691–211703.
- [175] Xuren Wang, Runshi Liu, Jie Yang, Rong Chen, Zhiting Ling, Peian Yang, Kai Zhang, Cyber threat intelligence entity extraction based on deep learning and field knowledge engineering, in: 2022 IEEE 25th International Conference on Computer Supported Cooperative Work in Design, *CSCWD*, IEEE, 2022, pp. 406–413.
- [176] Yongyan Guo, Zhengyu Liu, Cheng Huang, Jiayong Liu, Wangyuan Jing, Ziwang Wang, Yanghao Wang, CyberRel: Joint entity and relation extraction for cybersecurity concepts, in: Information and Communications Security: 23rd International Conference, *ICICS 2021*, Chongqing, China, November 19–21, 2021, Proceedings, Part I 23, Springer, 2021, pp. 447–463.
- [177] Jun Zhao, Qiben Yan, Xudong Liu, Bo Li, Guangsheng Zuo, Cyber threat intelligence modeling based on heterogeneous graph convolutional network, in: 23rd International Symposium on Research in Attacks, Intrusions and Defenses, *RAID 2020*, 2020, pp. 241–256.
- [178] Yitong Ren, Yanjun Xiao, Yinghai Zhou, Zhiyong Zhang, Zhihong Tian, CSKG4APT: A cybersecurity knowledge graph for advanced persistent threat organization attribution, *IEEE Trans. Knowl. Data Eng.* (2022).
- [179] Chen Gao, Xuan Zhang, Hui Liu, Data and knowledge-driven named entity recognition for cyber security, *Cybersecurity* 4 (1) (2021) 1–13.
- [180] Ning Luo, Xiangyu Du, Yitong He, Jun Jiang, Xuren Wang, Zhengwei Jiang, Kai Zhang, A framework for document-level cybersecurity event extraction from open source data, in: 2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design, *CSCWD*, IEEE, 2021, pp. 422–427.
- [181] Xuren Wang, Songheng He, Zihan Xiong, Xinxin Wei, Zhengwei Jiang, Sihao Chen, Jun Jiang, APTNER: A specific dataset for NER missions in cyber threat intelligence field, in: 2022 IEEE 25th International Conference on Computer Supported Cooperative Work in Design, *CSCWD*, IEEE, 2022, pp. 1233–1238.
- [182] Dongying Yang, Tao Lian, Wen Zheng, Cai Zhao, Enriching word information representation for Chinese cybersecurity named entity recognition, *Neural Process. Lett.* (2023) 1–19.
- [183] Ying Wei, Lili Bo, Xiaobing Sun, Bin Li, Tao Zhang, Chuanqi Tao, Automated event extraction of CVE descriptions, *Inf. Softw. Technol.* 158 (2023) 107178.
- [184] Pontus Stenetorp, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, Jun'ichi Tsujii, BRAT: a web-based tool for NLP-assisted text annotation, in: Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics, 2012, pp. 102–107.
- [185] Explosion, Prodigy, 2017, <https://prodigy.ai/>.
- [186] Alexandros Zacharis, Constantinos Patsakis, AiCEF: an AI-assisted cyber exercise content generation framework using named entity recognition, *Int. J. Inf. Secur.* (2023) 1–22.
- [187] Jie Yang, Yue Zhang, Linwei Li, Xingxuan Li, YEDDA: A lightweight collaborative text span annotation tool, 2017, arXiv preprint arXiv:1711.03759.
- [188] Jing Li, Aixin Sun, Jianglei Han, Chenliang Li, A survey on deep learning for named entity recognition, *IEEE Trans. Knowl. Data Eng.* 34 (1) (2020) 50–70.
- [189] Pingchuan Ma, Bo Jiang, Zhigang Lu, Ning Li, Zhengwei Jiang, Cybersecurity named entity recognition using bidirectional long short-term memory with conditional random fields, *Tsinghua Sci. Technol.* 26 (3) (2020) 259–265.
- [190] Housseem Gasmı, Abdelaziz Bouras, Jannik Laval, Encoder recurrent neural networks for cybersecurity named entity recognition, *ICSEA 11* (2018) 2018.
- [191] Tiberiu-Marian Georgescu, Bogdan Iancu, Madalina Zuerini, Named-entity-recognition-based automated system for diagnosing cybersecurity situations in IoT networks, *Sensors* 19 (15) (2019) 3380.
- [192] Ya Qin, Guo-wei Shen, Wen-bo Zhao, Yan-ping Chen, Miao Yu, Xin Jin, A network security entity recognition method based on feature template and CNN-BiLSTM-CRF, *Front. Inf. Technol. Electron. Eng.* 20 (6) (2019) 872–884.
- [193] Arpita Roy, Youngja Park, SHimei Pan, Learning domain-specific word embeddings from sparse cybersecurity texts, 2017, arXiv preprint arXiv:1709.07470.
- [194] Yiming Cui, Ziqing Yang, Ting Liu, PERT: pre-training BERT with permuted language model, 2022, arXiv preprint arXiv:2203.06906.
- [195] Inoussa Mouiche, Sherif Saad, Entity and relation extractions for threat intelligence knowledge graphs, *Comput. Secur.* 148 (2025) 104120.
- [196] Robert A. Bridges, Corinne L. Jones, Michael D. Iannacone, Kelly M. Testa, John R. Goodall, Automatic labeling for entity extraction in cyber security, 2013, arXiv preprint arXiv:1308.4941.
- [197] Peipei Liu, Hong Li, Zuoguang Wang, Jie Liu, Yimo Ren, Hongsong Zhu, Multi-features based semantic augmentation networks for named entity recognition in threat intelligence, in: 2022 26th International Conference on Pattern Recognition, *ICPR*, IEEE, 2022, pp. 1557–1563.
- [198] Han Zhang, Yuanbo Guo, Tao Li, et al., Multifeature named entity recognition in information security based on adversarial learning, *Secur. Commun. Netw.* 2019 (2019).
- [199] Tao Li, Yongjin Hu, Ankang Ju, Zhuoran Hu, Adversarial active learning for named entity recognition in cybersecurity, *Comput. Mater. Contin.* 66 (1) (2021).
- [200] Ravendar Lal, et al., Information Extraction of Security Related Entities and Concepts from Unstructured Text, Faculty of the Graduate School of the University of Maryland, 2013.
- [201] Xuren Wang, Xinpei Liu, Shengqin Ao, Ning Li, Zhengwei Jiang, Zongyi Xu, Zihan Xiong, Mengbo Xiong, Xiaoqing Zhang, Dnrti: A large-scale dataset for named entity recognition in threat intelligence, in: 2020 IEEE 19th International Conference on Trust, Security and Privacy in Computing and Communications, *TrustCom*, IEEE, 2020, pp. 1842–1848.
- [202] Qian Li, Shu Guo, Jia Wu, Jianxin Li, Jiawei Sheng, Hao Peng, Lihong Wang, Event extraction by associating event types and argument roles, *IEEE Trans. Big Data* (2023).

- [203] Hieu Man Duc Trong, Duc-Trong Le, Amir Pouran Ben Veyseh, Thuat Nguyen, Thien Huu Nguyen, Introducing a new dataset for event detection in cyber-security texts, in: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP, 2020, pp. 5381–5390.
- [204] Rada Mihalcea, Paul Tarau, Textrank: Bringing order into text, in: Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing, 2004, pp. 404–411.
- [205] Mengmeng Tang, Yuanbo Guo, Qingchun Bai, Han Zhang, Trigger-free cyber-security event detection based on contrastive learning, *J. Supercomput.* (2023) 1–24.
- [206] Hyejin Shin, WooChul Shim, Jiin Moon, Jae Woo Seo, Sol Lee, Yong Ho Hwang, Cybersecurity event detection with new and re-emerging words, in: Proceedings of the 15th ACM Asia Conference on Computer and Communications Security, 2020, pp. 665–678.
- [207] Farah Mohammad, Saad Al-Ahmadi, Jalal Al-Muhtadi, Deep learning based cyber event detection from open-source re-emerging social data, *Comput. Mater. Contin.* 76 (2) (2023).
- [208] Taneeya Satyapanich, Tim Finin, Francis Ferraro, Extracting rich semantic information about cybersecurity events, in: 2019 IEEE International Conference on Big Data, Big Data, IEEE, 2019, pp. 5034–5042.
- [209] Jia Liu, Runzi Zhang, Wenmao Liu, Yinghua Zhang, Dujuan Gu, Mingkai Tong, Xingkai Wang, Jianxin Xue, Huanran Wang, Context2Vector: Accelerating security event triage via context representation learning, *Inf. Softw. Technol.* 146 (2022) 106856.
- [210] Ian Perera, Jena Hwang, Kevin Bayas, Bonnie Dorr, Yorick Wilks, Cyberattack prediction through public text analysis and mini-theories, in: 2018 IEEE International Conference on Big Data, Big Data, IEEE, 2018, pp. 3001–3010.
- [211] Wenhan Ge, Junfeng Wang, SeqMask: Behavior extraction over cyber threat intelligence via multi-instance learning, *Comput. J.* (2022) bxac172.
- [212] Chenjing Liu, Junfeng Wang, Xiangru Chen, Threat intelligence ATT&CK extraction based on the attention transformer hierarchical recurrent neural network, *Appl. Soft Comput.* 122 (2022) 108826.
- [213] Ehsan Aghaei, Ehab Al-Shaer, CVE-driven attack technique prediction with semantic information extraction and a domain-specific language model, 2023, arXiv preprint arXiv:2309.02785.
- [214] Zhiqiang Hao, Chuanyi Li, Xiao Fu, Bin Luo, Xiaojiang Du, Leveraging hierarchies: HMCAT for efficiently mapping CTI to attack techniques, in: European Symposium on Research in Computer Security, Springer, 2024, pp. 65–85.
- [215] Tianyi Wang, Kam Pui Chow, Automatic tagging of cyber threat intelligence unstructured data using semantics extraction, in: 2019 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2019, pp. 197–199.
- [216] Chia-Mei Chen, Fang-Hsuan Hsu, Jenq-Neng Hwang, Useful cyber threat intelligence relation retrieval using transfer learning, in: Proceedings of the 2023 European Interdisciplinary Cybersecurity Conference, 2023, pp. 42–46.
- [217] Xuren Wang, Rong Chen, Binghua Song, Jie Yang, Zhengwei Jiang, Xiaoqing Zhang, Xiaomeng Li, Shengqin Ao, A method for extracting unstructured threat intelligence based on dictionary template and reinforcement learning, in: 2021 IEEE 24th International Conference on Computer Supported Cooperative Work in Design, CSCWD, IEEE, 2021, pp. 262–267.
- [218] Wenli Shang, Bowen Wang, Pengcheng Zhu, Lei Ding, Shuang Wang, A span-based multivariate information-aware embedding network for joint relational triplet extraction of threat intelligence, *Knowl.-Based Syst.* 295 (2024) 111829.
- [219] Yang Liu, Xiaohui Han, Wenbo Zuo, Haiqing Lv, Jing Guo, CTI-JE: A joint extraction framework of entities and relations in unstructured cyber threat intelligence, in: 2024 27th International Conference on Computer Supported Cooperative Work in Design, CSCWD, IEEE, 2024, pp. 2728–2733.
- [220] Kashan Ahmed, Syed Khaldoon Khurshid, Sadaf Hina, CyberEntRel: Joint extraction of cyber entities and relations using deep learning, *Comput. Secur.* 136 (2024) 103579.
- [221] Xiaodi Wang, Zhonglin Liu, Jiayong Liu, Joint relational triple extraction with enhanced representation and binary tagging framework in cybersecurity, *Comput. Secur.* 144 (2024) 104001.
- [222] Haiqing Lv, Xiaohui Han, Hui Cui, Peipei Wang, Wenbo Zuo, Yang Zhou, Joint extraction of entities and relationships from cyber threat intelligence based on task-specific Fourier network, in: 2024 International Joint Conference on Neural Networks, IJCNN, IEEE, 2024, pp. 1–8.
- [223] Sofonias Yitagesu, Zhenchang Xing, Xiaowang Zhang, Zhiyong Feng, Xiaohong Li, Linyi Han, Unsupervised labeling and extraction of phrase-based concepts in vulnerability descriptions, in: 2021 36th IEEE/ACM International Conference on Automated Software Engineering, ASE, IEEE, 2021, pp. 943–954.
- [224] Erik Hemberg, Ashwin Srinivasan, Nick Rutar, Una-May O'Reilly, Sourcing language models and text information for inferring cyber threat, vulnerability and mitigation relationships, in: KDD, 2022.
- [225] Cainã Figueiredo Pereira, João Gabriel Lopes de Oliveira, Rodrigo Azevedo Santos, Daniel Vieira, Lucas Miranda, Gerson Zaverucha, Leandro Pfleger de Aguiar, Daniel Sadoc Menasché, A statistical relational learning approach towards products, software vulnerabilities and exploits, *IEEE Trans. Netw. Serv. Manag.* 20 (3) (2023) 3782–3802.
- [226] R. Vaitheeshwari, Eric Hsiao-Kuang Wu, Ying-Dar Lin, Ren-Hung Hwang, Po-Ching Lin, Yuan-Cheng Lai, Asad Ali, TRACE: Relationship analysis and causal factor extraction in cyber threat intelligence reports, *IEEE Trans. Dependable Secur. Comput.* (2025).
- [227] Ghaith Husari, Xi Niu, Bill Chu, Ehab Al-Shaer, Using entropy and mutual information to extract threat actions from cyber threat intelligence, in: 2018 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2018, pp. 1–6.
- [228] Xiu Ma, Leiqi Wang, Qiuqian Lv, Yan Wang, Qi Zhang, Jianguo Jiang, CyEvent2vec: Attributed heterogeneous information network based event embedding framework for cyber security events analysis, in: 2022 International Joint Conference on Neural Networks, IJCNN, IEEE, 2022, pp. 01–08.
- [229] Huixia Zhang, Guowei Shen, Chun Guo, Yunhe Cui, Chaohui Jiang, Ex-action: Automatically extracting threat actions from cyber threat intelligence report based on multimodal learning, *Secur. Commun. Netw.* 2021 (2021) 1–12.
- [230] Yi-Ting Huang, Chi Yu Lin, Ying-Ren Guo, Kai-Chieh Lo, Yeali S. Sun, Meng Chang Chen, Open source intelligence for malicious behavior discovery and interpretation, *IEEE Trans. Dependable Secur. Comput.* 19 (2) (2021) 776–789.
- [231] Ming Fan, Xiapu Luo, Jun Liu, Chunyin Nong, Qinghua Zheng, Ting Liu, Ctdroid: leveraging a corpus of technical blogs for android malware analysis, *IEEE Trans. Reliab.* 69 (1) (2019) 124–138.
- [232] Mahender Kumar, Gregory Epiphaniou, Carsten Maple, Leveraging semantic relationships to prioritise indicators of compromise in additive manufacturing systems, 2023, arXiv preprint arXiv:2305.04102.
- [233] Amirreza Niakanlahiji, Jinpeng Wei, Bei-Tseng Chu, A natural language processing based trend analysis of advanced persistent threat techniques, in: 2018 IEEE International Conference on Big Data, Big Data, IEEE, 2018, pp. 2995–3000.
- [234] Luis Felipe Gutierrez, Akbar Namin, Contextminer: Mining contextual features for conceptualizing knowledge in security texts, *IEEE Access* 10 (2022) 85891–85904.
- [235] Matthew Sills, Priyanka Ranade, Sudip Mittal, Cybersecurity threat intelligence augmentation and embedding improvement-a healthcare usecase, in: 2020 IEEE International Conference on Intelligence and Security Informatics, ISI, IEEE, 2020, pp. 1–6.
- [236] F. Böhm, F. Menges, G. Pernul, Graph-based visual analytics for cyber threat intelligence, *cybersecurity* 1 (1): 1–19, 2018.
- [237] Sudip Mittal, Anupam Joshi, Tim Finin, Cyber-all-intel: An AI for security related threat intelligence, 2019, arXiv preprint arXiv:1905.02895.
- [238] Md Tanvirul Alam, Dipkamal Bhusal, Youngja Park, Nidhi Rastogi, Looking beyond IoCs: Automatically extracting attack patterns from external CTI, 2022, arXiv preprint arXiv:2211.01753.
- [239] Aditya Pingle, Aritran Piplai, Sudip Mittal, Anupam Joshi, James Holt, Richard Zak, Relext: Relation extraction using deep learning approaches for cybersecurity knowledge graph improvement, in: Proceedings of the 2019 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining, 2019, pp. 879–886.
- [240] Yali Gao, Xiaoyong Li, Hao Peng, Binxing Fang, S. Yu Philip, Hincti: A cyber threat intelligence modeling and identification system based on heterogeneous information network, *IEEE Trans. Knowl. Data Eng.* 34 (2) (2020) 708–722.
- [241] Z. Li, J. Zeng, Y. Chen, Z. Liang, AttackKG: Constructing technique knowledge graph from cyber threat intelligence reports", 2022, arXiv. (Accessed 02 August 2022).
- [242] Junbo Jia, Li Yang, Yuchen Wang, Anyuan Sang, Hyper attack graph: Constructing a hypergraph for cyber threat intelligence analysis, *Comput. Secur.* 149 (2025) 104194.
- [243] Gaosheng Wang, Peipei Liu, Jintao Huang, Haoyu Bin, Xi Wang, Hongsong Zhu, KnowCTI: Knowledge-based cyber threat intelligence entity and relation extraction, *Comput. Secur.* 141 (2024) 103824.
- [244] Zong-Xun Li, Yu-Jun Li, Yi-Wei Liu, Cheng Liu, Nan-Xin Zhou, K-CTIAA: Automatic analysis of cyber threat intelligence based on a knowledge graph, *Symmetry* 15 (2) (2023) 337.
- [245] Chunyan Ma, Zhengwei Jiang, Kai Zhang, Zhiting Ling, Jun Jiang, Yizhe You, Peian Yang, Huamin Feng, TIMFuser: A multi-granular fusion framework for cyber threat intelligence, *Comput. Secur.* 148 (2025) 104141.
- [246] Florian Klaus Kaiser, Uriel Dardik, Aviad Elizur, Polina Zilberman, Nir Daniel, Marcus Wiens, Frank Schultmann, Yuval Elovici, Rami Puzis, Attack hypotheses generation based on threat intelligence knowledge graph, *IEEE Trans. Dependable Secur. Comput.* (2023).
- [247] Hyeonseong Jo, Yongjae Lee, Seungwon Shin, Vulcan: Automatic extraction and analysis of cyber threat intelligence from unstructured text, *Comput. Secur.* 120 (2022) 102763.
- [248] Yongyan Guo, Zhengyu Liu, Cheng Huang, Nannan Wang, Hai Min, Wenbo Guo, Jiayong Liu, A framework for threat intelligence extraction and fusion, *Comput. Secur.* 132 (2023) 103371.
- [249] Zhenpeng Shi, Nikolay Matyunin, Kalman Graffi, David Starobinski, Uncovering cwe-cve-cpe relations with threat knowledge graphs, *ACM Trans. Priv. Secur.* 27 (1) (2024) 1–26.
- [250] Ezekia Gilliard, Jinshuo Liu, Ahmed Abubakar Aliyu, Knowledge graph reasoning for cyber attack detection, *IET Commun.* 18 (4) (2024) 297–308.

- [251] Saad El Jaouhari, Shaikh Ahmed, CTIoT: A cyber threat intelligence tool for IoT, in: 2024 International Wireless Communications and Mobile Computing, IWCMC, IEEE, 2024, pp. 1483–1489.
- [252] Yuta Kazato, Yoshihide Nakagawa, Yuichi Nakatani, Improving maliciousness estimation of indicator of compromise using graph convolutional networks, in: 2020 IEEE 17th Annual Consumer Communications & Networking Conference, CCNC, IEEE, 2020, pp. 1–7.
- [253] Siwar Kriaa, Yahia Chaabane, SeckKG: Leveraging attack detection and prediction using knowledge graphs, in: 2021 12th International Conference on Information and Communication Systems, ICICS, IEEE, 2021, pp. 112–119.
- [254] Jun Zhao, Minglai Shao, Hong Wang, Xiaomei Yu, Bo Li, Xudong Liu, Cyber threat prediction using dynamic heterogeneous graph learning, *Knowl.-Based Syst.* 240 (2022) 108086.
- [255] Yang Zhang, Jiarui Chen, Zhe Cheng, Xiong Shen, Jiancheng Qin, Yingzheng Han, Yiqin Lu, Edge propagation for link prediction in requirement-cyber threat intelligence knowledge graph, *Inform. Sci.* 653 (2024) 119770.
- [256] Elmar Kiesling, Andreas Ekelhart, Kabul Kurniawan, Fajar Ekaputra, The SEPSES knowledge graph: An integrated resource for cybersecurity, in: International Semantic Web Conference, Springer, 2019, pp. 198–214.
- [257] Jun Zhao, Xudong Liu, Qiben Yan, Bo Li, Minglai Shao, Hao Peng, Lichao Sun, Automatically predicting cyber attack preference with attributed heterogeneous attention networks and transductive learning, *Comput. Secur.* 102 (2021) 102152.
- [258] Burak Gulbay, Mehmet Demirci, APT-scope: A novel framework to predict advanced persistent threat groups from enriched heterogeneous information network of cyber threat intelligence, *Eng. Sci. Technol. Int. J.* 57 (2024) 101791.
- [259] Yongfei Li, Yuanbo Guo, Chen Fang, Yingze Liu, Qingli Chen, et al., A novel threat intelligence information extraction system combining multiple models, *Secur. Commun. Netw.* 2022 (2022).
- [260] Soham Dasgupta, Aritran Piplai, Priyanka Ranade, Anupam Joshi, Cybersecurity knowledge graph improvement with graph neural networks, in: 2021 IEEE International Conference on Big Data, Big Data, IEEE, 2021, pp. 3290–3297.
- [261] Shaswata Mitra, Aritran Piplai, Sudip Mittal, Anupam Joshi, Combating fake cyber threat intelligence using provenance in cybersecurity knowledge graphs, in: 2021 IEEE International Conference on Big Data, Big Data, IEEE, 2021, pp. 3316–3323.
- [262] Gautam Srivastava, Rutvij H. Jhaveri, Sweta Bhattacharya, Sharnil Pandya, Praveen Kumar Reddy Maddikunta, Gokul Yenduri, Jon G. Hall, Mamoun Alazab, Thippa Reddy Gadekallu, et al., XAI for cybersecurity: state of the art, challenges, open issues and future directions, 2022, arXiv preprint arXiv: 2206.03585.
- [263] Zhibo Zhang, Hussam Al Hamadi, Ernesto Damiani, Chan Yeob Yeun, Fatma Taher, Explainable artificial intelligence applications in cyber security: State-of-the-art in research, *IEEE Access* (2022).
- [264] Harsha Moraliyage, Vidura Sumanasena, Daswin De Silva, Rashmika Nawaratne, Lina Sun, Daminda Alahakoon, Multimodal classification of onion services for proactive cyber threat intelligence using explainable deep learning, *IEEE Access* 10 (2022) 56044–56056.
- [265] Huiling Li, Jun Wu, Hansong Xu, Gaolei Li, Mohsen Guizani, Explainable intelligence-driven defense mechanism against advanced persistent threats: A joint edge game and AI approach, *IEEE Trans. Dependable Secur. Comput.* 19 (2) (2021) 757–775.
- [266] Qingqin Wang, Hanbing Yan, Zhihui Han, Explainable apt attribution for malware using nlp techniques, in: 2021 IEEE 21st International Conference on Software Quality, Reliability and Security, QRS, IEEE, 2021, pp. 70–80.
- [267] Wenhan Ge, Junfeng Wang, Tongcan Lin, Binhui Tang, Xiaohui Li, Explainable cyber threat behavior identification based on self-adversarial topic generation, *Comput. Secur.* 132 (2023) 103369.
- [268] Cynthia Wagner, Alexandre Dulaunoy, Gérard Wagnen, Andras Iklody, Misp: The design and implementation of a collaborative threat intelligence sharing platform, in: Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security, 2016, pp. 49–56.
- [269] Borce Stojkovski, Gabriele Lenzi, Vincent Koenig, Salvador Rivas, What's in a cyber threat intelligence sharing platform? A mixed-methods user experience investigation of MISP, in: Annual Computer Security Applications Conference, 2021, pp. 385–398.
- [270] Han Wang, Alfonso Iacovazzi, Seonghyun Kim, Shahid Raza, CLEVER: Crafting intelligent MISP for cyber threat intelligence, in: 2024 IEEE 49th Conference on Local Computer Networks, LCN, IEEE, 2024, pp. 1–9.
- [271] Hyung-Woo Lee, Malware attack indicators of compromise generation and sharing mechanism for cyber threat intelligence, in: Proceedings of the 2024 13th International Conference on Software and Computer Applications, 2024, pp. 90–96.
- [272] Valeria Valdés Ríos, Fatiha Zaidi, Ana Rosa Cavalli, Angel Rego, Towards the adoption of automated cyber threat intelligence information sharing with integrated risk assessment, in: Proceedings of the 19th International Conference on Availability, Reliability and Security, 2024, pp. 1–9.
- [273] Eric Jollès, Sébastien Gillard, Dimitri Percia David, Martin Strohmeier, Alain Mermoud, Building collaborative cybersecurity for critical infrastructure protection: Empirical evidence of collective intelligence information sharing dynamics on ThreatFox, in: International Conference on Critical Information Infrastructures Security, Springer, 2022, pp. 140–157.
- [274] Alessandra de Melo e Silva, João José Costa Gondim, Robson de Oliveira Albuquerque, Luis Javier García Villalba, A methodology to evaluate standards and platforms within cyber threat intelligence, *Futur. Internet* 12 (6) (2020) 108.
- [275] Thomas D. Wagner, Esther Palomar, Khaled Mahbub, Ali E. Abdallah, Towards an anonymity supported platform for shared cyber threat intelligence, in: Risks and Security of Internet and Systems: 12th International Conference, CRISIS 2017, Dinard, France, September 19–21, 2017, Revised Selected Papers 12, Springer, 2018, pp. 175–183.
- [276] Md Farhan Haque, Ram Krishnan, Toward automated cyber defense with secure sharing of structured cyber threat intelligence, *Inf. Syst. Front.* (2021) 1–14.
- [277] Gustavo González-Granadillo, Mario Faiella, Ibéria Medeiros, Rui Azevedo, Susana González-Zarzosa, ETIP: An enriched threat intelligence platform for improving OSINT correlation, analysis, visualization and sharing capabilities, *J. Inf. Secur. Appl.* 58 (2021) 102715.
- [278] Florian Menges, Benedikt Putz, Günther Pernul, DEALER: decentralized incentives for threat intelligence reporting and exchange, *Int. J. Inf. Secur.* 20 (5) (2021) 741–761.
- [279] Davy Preuveneers, Wouter Joosen, Jorge Bernal Bernabe, Antonio Skarmeta, Distributed security framework for reliable threat intelligence sharing, *Secur. Commun. Netw.* 2020 (2020).
- [280] Davy Preuveneers, Wouter Joosen, TATIS: trustworthy APIs for threat intelligence sharing with UMA and CP-ABE, in: Foundations and Practice of Security: 12th International Symposium, FPS 2019, Toulouse, France, November 5–7, 2019, Revised Selected Papers 12, Springer, 2020, pp. 172–188.
- [281] Sean Barnum, Standardizing cyber threat intelligence information with the structured threat information expression (stix), *Mitre Corp.* 11 (2012) 1–22.
- [282] Raúl Riesco, Víctor A. Villagrà, Leveraging cyber threat intelligence for a dynamic risk framework: Automation by using a semantic reasoner and a new combination of standards (STIX™, SWRL and OWL), *Int. J. Inf. Secur.* 18 (6) (2019) 715–739.
- [283] Florian Menges, Günther Pernul, A comparative analysis of incident reporting formats, *Comput. Secur.* 73 (2018) 87–101.
- [284] Andrew Ramsdale, Stavros Shialeles, Nicholas Kolokotronis, A comparative analysis of cyber-threat intelligence sources, formats and languages, *Electronics* 9 (5) (2020) 824.
- [285] Vitor Jesus, Balraj Bains, Victor Chang, Sharing is caring: Hurdles and prospects of open, crowd-sourced cyber threat intelligence, *IEEE Trans. Eng. Manage.* (2023).
- [286] Christian Sillaber, Clemens Sauerwein, Andrea Mussmann, Ruth Brey, Data quality challenges and future research directions in threat intelligence sharing practice, in: Proceedings of the 2016 ACM on Workshop on Information Sharing and Collaborative Security, 2016, pp. 65–70.
- [287] Thomas Geras, Thomas Schreck, The “big beast to tackle”: Practices in quality assurance for cyber threat intelligence, in: Proceedings of the 27th International Symposium on Research in Attacks, Intrusions and Defenses, 2024, pp. 337–352.
- [288] Daniel Schlette, Fabian Böhm, Marco Caselli, Günther Pernul, Measuring and visualizing cyber threat intelligence quality, *Int. J. Inf. Secur.* 20 (2021) 21–38.
- [289] Adam Zibak, Clemens Sauerwein, Andrew C. Simpson, Threat intelligence quality dimensions for research and practice, *Digit. Threat.: Res. Pr.* 3 (4) (2022) 1–22.
- [290] Li Qiang, Jiang Zhengwei, Yang Zeming, Liu Baoxu, Wang Xin, Zhang Yunan, A quality evaluation method of cyber threat intelligence in user perspective, in: 2018 17th IEEE International Conference on Trust, Security and Privacy in Computing and Communications/12th IEEE International Conference on Big Data Science and Engineering, TrustCom/BigDataSE, IEEE, 2018, pp. 269–276.
- [291] Thomas Schaberreiter, Veronika Kupfersberger, Konstantinos Rantos, Arnolnt Spyros, Alexandros Papanikolaou, Christos Ilioudis, Gerald Quirchmayr, A quantitative evaluation of trust in the quality of cyber threat intelligence sources, in: Proceedings of the 14th International Conference on Availability, Reliability and Security, 2019, pp. 1–10.
- [292] Harm Griffioen, Tim Booij, Christian Doerr, Quality evaluation of cyber threat intelligence feeds, in: Applied Cryptography and Network Security: 18th International Conference, ACNS 2020, Rome, Italy, October 19–22, 2020, Proceedings, Part II 18, Springer, 2020, pp. 277–296.
- [293] Sheng-Shan Chen, Ren-Hung Hwang, Asad Ali, Ying-Dar Lin, Yu-Chih Wei, Tun-Wen Pai, Improving quality of indicators of compromise using STIX graphs, *Comput. Secur.* 144 (2024) 103972.
- [294] Beomjin Jin, Eunsoo Kim, Hyunwoo Lee, Elisa Bertino, Doowon Kim, Hyoungshick Kim, Sharing cyber threat intelligence: Does it really help? in: Proceedings of the 31st Annual Network and Distributed System Security Symposium, NDSS, 2024.
- [295] Hongyi Li, Ze Shi, Chengwei Pan, Di Zhao, Nan Sun, Cybersecurity knowledge graphs construction and quality assessment, *Complex Intell. Syst.* 10 (1) (2024) 1201–1217.

- [296] Libin Yang, Menghan Wang, Wei Lou, An automated dynamic quality assessment method for cyber threat intelligence, *Comput. Secur.* 148 (2025) 104079.
- [297] Thomas D. Wagner, Esther Palomar, Khaled Mahbub, Ali E. Abdallah, A novel trust taxonomy for shared cyber threat intelligence, *Secur. Commun. Netw.* 2018 (2018).
- [298] David W. Chadwick, Wenjun Fan, Gianpiero Costantino, Rogério De Lemos, Francesco Di Cerbo, Ian Herwono, Mirko Manea, Paolo Mori, Ali Sajjad, Xiao-Si Wang, A cloud-edge based data security architecture for sharing and analysing cyber threat information, *Future Gener. Comput. Syst.* 102 (2020) 710–722.
- [299] Tongtong Jiang, Guowei Shen, Chun Guo, Yunhe Cui, Bo Xie, BFLS: Blockchain and federated learning for sharing threat detection models as cyber threat intelligence, *Comput. Netw.* 224 (2023) 109604.
- [300] Daire Homan, Ian Shiel, Christina Thorpe, A new network model for cyber threat intelligence sharing using blockchain technology, in: 2019 10th IFIP International Conference on New Technologies, Mobility and Security, NTMS, IEEE, 2019, pp. 1–6.
- [301] Soumya Purohit, Roshan Neupane, Naga Ramya Bhamidipati, Varsha Vakka-vanthula, Songjie Wang, Matthew Rockey, Prasad Calyam, Cyber threat intelligence sharing for co-operative defense in multi-domain entities, *IEEE Trans. Dependable Secur. Comput.* (2022).
- [302] Raúl Riesco, Xavier Larriva-Novo, Víctor A. Villagrà, Cybersecurity threat intelligence knowledge exchange based on blockchain: Proposal of a new incentive model based on blockchain and smart contracts to foster the cyber threat and risk intelligence exchange of information, *Telecommun. Syst.* 73 (2) (2020) 259–288.
- [303] Philip Huff, Qinghua Li, A distributed ledger for non-attributable cyber threat intelligence exchange, in: *Security and Privacy in Communication Networks: 17th EAI International Conference, SecureComm 2021, Virtual Event, September 6–9, 2021, Proceedings, Part I 17*, Springer, 2021, pp. 164–184.
- [304] Tian Li, Anit Kumar Sahu, Ameet Talwalkar, Virginia Smith, Federated learning: Challenges, methods, and future directions, *IEEE Signal Process. Mag.* 37 (3) (2020) 50–60.
- [305] Dincy R. Arikkat, Mert Cihangiroglu, Mauro Conti, K.A. Rafidha Rehiman, Serena Nicolazzo, Antonino Nocera, P. Vinod, SeCTIS: A framework to secure CTI sharing, *Future Gener. Comput. Syst.* 164 (2025) 107562.
- [306] Saifur Rahman, Shantanu Pal, Zahra Jadidi, Chandan Karmakar, Robust cyber threat intelligence sharing using federated learning for smart grids, *IEEE Trans. Comput. Soc. Syst.* (2024).
- [307] Kealan Dunnett, Shantanu Pal, Zahra Jadidi, Volkan Dedeoglu, Raja Jurdak, Priv-Share: A privacy-preserving framework for differential and trustless delegation of cyber threat intelligence using blockchain, *Comput. Netw.* 252 (2024) 110686.
- [308] Nicolas Papernot, Patrick McDaniel, Ananthram Swami, Richard Harang, Crafting adversarial input sequences for recurrent neural networks, in: *MILCOM 2016-2016 IEEE Military Communications Conference*, IEEE, 2016, pp. 49–54.
- [309] Anay Pattanaik, Zhenyi Tang, Shuijing Liu, Gautham Bommanna, Girish Chowdhary, Robust deep reinforcement learning with adversarial attacks, 2017, arXiv preprint arXiv:1712.03632.
- [310] Shilin Qiu, Qihui Liu, Shijie Zhou, Wen Huang, Adversarial attack and defense technologies in natural language processing: A survey, *Neurocomputing* 492 (2022) 278–307.
- [311] Basemah Alshemali, Jugal Kalita, Improving the reliability of deep neural networks in NLP: A review, *Knowl.-Based Syst.* 191 (2020) 105210.
- [312] Wei Emma Zhang, Quan Z. Sheng, Ahoud Alhazmi, Chenliang Li, Adversarial attacks on deep-learning models in natural language processing: A survey, *ACM Trans. Intell. Syst. Technol. (TIST)* 11 (3) (2020) 1–41.
- [313] MITRE, Virus total data poisoning case studies, 2021, <https://github.com/mitre/advmlthreatmatrix/blob/master/pages/case-studies-page.md#virustotal-poisoning>.
- [314] Tian Jiang, Yunqi Liu, Xiaohui Cui, Textual adversarial attacks in cybersecurity named entity recognition, *Comput. Secur.* 150 (2025) 104278.
- [315] Ji Gao, Jack Lanchantin, Mary Lou Soffa, Yanjun Qi, Black-box generation of adversarial text sequences to evade deep learning classifiers, in: *2018 IEEE Security and Privacy Workshops, SPW, IEEE*, 2018, pp. 50–56.
- [316] Samson Tan, Shafiq Joty, Min-Yen Kan, Richard Socher, It's Morphin'time! combating linguistic discrimination with inflectional perturbations, 2020, arXiv preprint arXiv:2005.04364.
- [317] Di Jin, Zhijing Jin, Joey Tianyi Zhou, Peter Szolovits, Is bert really robust? a strong baseline for natural language attack on text classification and entailment, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8018–8025.
- [318] Alberto Maria Mongardini, Massimo La Morgia, Sushil Jajodia, Luigi Vincenzo Mancini, Alessandro Mei, DARD: Deceptive approaches for robust defense against IP theft, *IEEE Trans. Inf. Forensics Secur.* (2024).
- [319] Zhihua Li, Xinye Yu, Yukai Zhao, A web semantic mining method for fake cybersecurity threat intelligence in open source communities, *Int. J. Semant. Web Inf. Syst. (IJSWIS)* 20 (1) (2024) 1–22.
- [320] Nitika Khurana, Sudip Mittal, Aritran Piplai, Anupam Joshi, Preventing poisoning attacks on AI based threat intelligence systems, in: *2019 IEEE 29th International Workshop on Machine Learning for Signal Processing, MLSP, IEEE*, 2019, pp. 1–6.
- [321] Jiehui Liu, Jieyu Zhan, Constructing knowledge graph from cyber threat intelligence using large language model, in: *2023 IEEE International Conference on Big Data, BigData, IEEE*, 2023, pp. 516–521.
- [322] Yuelin Hu, Futai Zou, Jiajia Han, Xin Sun, Yilei Wang, Llm-tikg: Threat intelligence knowledge graph construction utilizing large language model, *Comput. Secur.* 145 (2024) 103999.
- [323] Romy Fieblinger, Md Tanvirul Alam, Nidhi Rastogi, Actionable cyber threat intelligence using knowledge graphs and large language models, in: *2024 IEEE European Symposium on Security and Privacy Workshops, EuroS&PW, IEEE*, 2024, pp. 100–111.
- [324] Yongheng Zhang, Tingwen Du, Yunshan Ma, Xiang Wang, Yi Xie, Guozheng Yang, Yuliang Lu, Ee-Chien Chang, AttacKG+: Boosting attack graph construction with large language models, *Comput. Secur.* 150 (2025) 104220.
- [325] Matteo Boffa, Idilio Drago, Marco Mellia, Luca Vassio, Danilo Giordano, Rodolfo Valentim, Zied Ben Houidi, LogPrécis: Unleashing language models for automated malicious log analysis: Précis: A concise summary of essential points, statements, or facts, *Comput. Secur.* 141 (2024) 103805.
- [326] Minghao Chen, Kaijie Zhu, Bin Lu, Ding Li, Qingjun Yuan, Yuefei Zhu, AECR: Automatic attack technique intelligence extraction based on fine-tuned large language model, *Comput. Secur.* 150 (2025) 104213.
- [327] Samaneh Shafee, Alysso Bessani, Pedro M. Ferreira, Evaluation of LLM-based chatbots for OSINT-based cyber threat awareness, *Expert Syst. Appl.* (2024) 125509.
- [328] Saad Ullah, Mingji Han, Saurabh Pujar, Hammond Pearce, Ayse Coskun, Gianluca Stringhini, LLMs cannot reliably identify and reason about security vulnerabilities (yet?): A comprehensive evaluation, framework, and benchmarks, in: *IEEE Symposium on Security and Privacy*, 2024.
- [329] Marcus Botacin, GPTthreats-3: Is automatic malware generation a threat? in: *2023 IEEE Security and Privacy Workshops, SPW*, 2023, pp. 238–254.
- [330] Khalifa Afane, Wenqi Wei, Ying Mao, Junaid Farooq, Juntao Chen, Next-generation phishing: How LLM agents empower cyber attackers, in: *2024 IEEE International Conference on Big Data, BigData, IEEE*, 2024, pp. 2558–2567.