



STIXnet: A Novel and Modular Solution for Extracting All STIX Objects in CTI Reports

Francesco Marchiori

University of Padova

Padua, Italy

francesco.marchiori@math.unipd.it

Mauro Conti

University of Padova

Padua, Italy

mauro.conti@unipd.it

Nino Vincenzo Verde

Leonardo S.p.A.

Rome, Italy

nino.verde@leonardo.com

ABSTRACT

The automatic extraction of information from Cyber Threat Intelligence (CTI) reports is crucial in risk management. The increased frequency of the publications of these reports has led researchers to develop new systems for automatically recovering different types of entities and relations from textual data. Most state-of-the-art models leverage Natural Language Processing (NLP) techniques, which perform greatly in extracting a few types of entities at a time but cannot detect heterogeneous data or their relations. Furthermore, several paradigms, such as STIX, have become de facto standards in the CTI community and dictate a formal categorization of different entities and relations to enable organizations to share data consistently.

This paper presents STIXnet, the first solution for the automated extraction of all STIX entities and relationships in CTI reports. Through the use of NLP techniques and an interactive Knowledge Base (KB) of entities, our approach obtains F1 scores comparable to state-of-the-art models for entity extraction (0.916) and relation extraction (0.724) while considering significantly more types of entities and relations. Moreover, STIXnet constitutes a modular and extensible framework that manages and coordinates different modules to merge their contributions uniquely and exhaustively. With our approach, researchers and organizations can extend their Information Extraction (IE) capabilities by integrating the efforts of several techniques without needing to develop new tools from scratch.

CCS CONCEPTS

• **Security and privacy**; • **Information systems** → *Information extraction*;

KEYWORDS

Cyber Threat Intelligence, Natural Language Processing, Information Extraction, STIX

ACM Reference Format:

Francesco Marchiori, Mauro Conti, and Nino Vincenzo Verde. 2023. STIXnet: A Novel and Modular Solution for Extracting All STIX Objects in CTI Reports. In *The 18th International Conference on Availability, Reliability and*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ARES 2023, August 29–September 01, 2023, Benevento, Italy

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0772-8/23/08...\$15.00

<https://doi.org/10.1145/3600160.3600182>

Security (ARES 2023), August 29–September 01, 2023, Benevento, Italy. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3600160.3600182>

1 INTRODUCTION

The increasing number of cyberattacks has concerned many people and organizations, which, now more than ever, are potentially exposing their data to breaches and attacks. In particular, companies must be aware of Advanced Persistent Threats (APTs), stealthy actors that establish a persistent presence in networks to steal information. APTs constitute one of the biggest concerns for most organizations since they are hard to identify and use complex attacks that are difficult to prevent [5]. Security experts use various types of intelligence to track their movements, motivations, and behavior to counter them. To do this, researchers analyze previous breaches to understand the techniques that APTs use, the vulnerabilities and malwares they exploit, and the tools used for their deployment [3].

The collection and distribution of these kinds of data fall in the Cyber Threat Intelligence (CTI) field. Many researchers are involved in this discipline since, through this intelligence, companies can proactively defend against specific threat actors that might target them. In this way, organizations can protect specific assets more at risk than others and redirect their security budget in the best possible way [23]. For these reasons, companies must be up-to-date on the latest attacks, malwares, and techniques. To do so, they must collect intelligence by analyzing previous incidents or using external sources. Indeed, the distribution of this intelligence is provided by different vendors through reports and bulletins or by Open Source Intelligence (OSINT, i.e., any data that can be gathered for free). These reports are usually written in English and contain all the information on a particular incident or actor [28].

The extraction of relevant information from these reports is usually performed by CTI analysts, trained to recognize the entities of interest and the relations between them. Once identified, intelligence must be annotated following certain standards to ensure machine readability and compatibility during distribution. One of the community's most popular and recognized standards is the Structured Threat Information Expression (STIX) language [2], categorizing intelligence into different entities and relations types. However, this action can be time-consuming, given the reports' length and increased publication frequency in the last few years. To solve this issue, many attempts have been made to extract entities automatically and relations [9, 13, 35]. However, most state-of-the-art models focus on extracting a single or a few types of intelligence at a time, in which they obtain excellent results [14, 31]. In most real-world scenarios, the extraction of only a subset of entities and

relations might not be enough to fully characterize the data contained in a report. Despite that, merging the results of different models is not trivial since conflicts might appear in the extracted data, and intelligence is often displayed through different standards and paradigms.

Contribution. In this paper, we present STIXnet, the first modular and extensible system for the automated extraction of all STIX entities and relationships in CTI reports. STIXnet works by leveraging different techniques, such as Natural Language Processing (NLP), to extract threat intelligence from the text of the report and identify the relevant pieces of information while also retrieving the relations among them. Our tool uses a rich Knowledge Base (KB) that contains CTI data from various sources and previous report extractions. The Knowledge Base can be enlarged with each execution and provide data for training Machine Learning (ML) and Deep Learning (DL) models used in some STIXnet modules. Through a graphical interface, the results of the STIXnet processing can be visualized as a graph in which nodes constitute entities and edges constitute relations. Each node can then be expanded with additional information stored in the database and thus provide a quick and interactive overview for each entity in the Knowledge Base. The architecture of our solution allows for the extraction of all entities and relations compliant with the STIX standard without running into structural constraints such as model retraining and dataset reannotation. In Table 1, we show in more detail the types of entities that STIXnet can extract compared to other models in the literature. It is worth noting that some of the models used for the comparison do not focus on STIX entities in particular, but their labels can be translated according to the STIX standard. Furthermore, in the comparison, we consider only entities that can be extracted from a report. Still, through our platform, it is possible to include entities compliant with the latest STIX 2.1 standard.

The main contributions of our work can be summarized as follows:

- We present **STIXnet**, the first system for automatically extracting all types of STIX entities (18) and relations (more than 100).
- We propose a novel framework for managing and coordinating several modules and their interaction for Information Extraction tasks. Through this framework, we propose a methodology for integrating results from different modules through a confidence value and with minimal supervision.
- We make our testbed (dataset and annotated reports through LabelStudio¹) and the code for some modules available at <https://github.com/Mhackiori/STIXnet/>.

Organization. The rest of the paper is organized as follows. In Section 2, we present some background on Natural Language Processing techniques and Cyber Threat Intelligence. Section 3 presents a review of previous works on Information Extraction (IE) from natural language reports in the field of CTI. The proposed methodology and the pipeline are presented in Section 4, followed by an analysis of the results in Section 5. Finally, Section 6 concludes this work.

2 BACKGROUND

We now give a more thorough background on the techniques that we use in the methodology (Section 2.1) and further expose the challenges of Cyber Threat Intelligence (Section 2.2).

2.1 Natural Language Processing

Given the fluency and convenience of natural language for human interactions, the field of Natural Language Processing is born to develop algorithms and models that can comprehend and analyze this type of language. These algorithms can also include Machine Learning or Deep Learning techniques, which create many opportunities for its applications. With ML and DL techniques, many documents can be automatically processed to extract named entities, detect their attributes, and retrieve their existing relations. For this reason, NLP has become particularly suitable for tasks involving the analysis of multiple reports and extracting a predefined set of data in a text [6].

NLP can be applied in every domain in which human language is the main vector through which information is conveyed, e.g., speech recognition (i.e., speech-to-text, the act of translating voice data into text data), Part-Of-Speech tagging (i.e., grammatical tagging or POS tagging, the act of determining the part of speech of a particular word based on its context) and Named Entity Recognition (i.e., NER, the act of identifying specific words in a text as a specific type of entity) [25]. In particular, this last technique is heavily used for Information Extraction tasks where a large number of text data is involved [18], such as medical applications [30], scientific research [15] and cybersecurity intelligence [35].

Information Extraction also comprises Relation Extraction, i.e., retrieving and classifying the semantic relationships between two (or more) tokens inside a text [12]. This task can become particularly important in sequence with the Entity Extraction task. In this way, the information in textual data becomes intertwined, and a knowledge graph of the processed text can be created [16].

2.2 Cyber Threat Intelligence

According to one of the definitions of the Computer Security Research Center at NIST², a cyber threat can be defined "*any circumstance or event with the potential to adversely impact organizational operations [...] via unauthorized access, destruction, disclosure, modification of information, and/or denial of service*". CTI is the field that studies these threats and analyzes the intentions of the threat actors, the techniques they use, and the tools and malwares they deploy. In this way, it is possible to profile the activities of the malicious actors and thus design more effective cyber defense strategies [28].

Several types of intelligence are treated by CTI, which can be more or less technical depending on the target audience that must digest it. Also, each piece of intel follows a specific life cycle from planning and direction to its dissemination and integration [19].

To efficiently share and distribute the collected information, a common protocol must be in place to avoid misunderstandings between the teams that consume it. To address these problems, the STIX (Structured Threat Information eXpression) standard has been created [2]. This standardized language has been created by

¹<https://labelstud.io/>

²https://csrc.nist.gov/glossary/term/cyber_threat

Table 1: Comparison of the types of entities extracted by STIXnet.

Model	Attack Pattern	Campaign	Course of Action	Identity	Indicator	Infrastructure	Intrusion Set	Location	Malware	Malware Analysis	Report	Threat Actor	Tool	Vulnerability
Weerawardhana et al. [31]					✓							✓	✓	✓
Li et al. [14]	✓								✓			✓	✓	
Zhou (2022) et al. [37]		✓		✓				✓				✓	✓	
Zhou (2023) et al. [36]		✓		✓				✓				✓	✓	
Ranade et al. [21]	✓		✓						✓			✓	✓	✓
Wang et al. [29]	✓			✓	✓			✓	✓			✓	✓	✓
STIXnet	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

MITRE³ and is driven by the collaboration of many individuals who keep it up-to-date. Including different types of entities and relationships makes it possible to accurately represent the information in a cyber security report through a STIX file. The latest version of the software is STIX 2.1⁴, which includes 18 STIX Domain Objects (SDOs) and more than 100 possible relations among them.

While STIX provides a standardized language for disseminating CTI, it does not provide intelligence. One of the most popular frameworks for Cyber Threat Intelligence is the MITRE ATT&CK framework, a publicly accessible Knowledge Base of TTPs extracted from real-world CTI reports that can be used as a foundation for building a personalized database of threat intelligence [24]. Entities in the ATT&CK Knowledge Base are categorized with different labels, of which the ones of interest for the scope of this paper are *tactics*, *techniques*, *groups*, and *software*. However, since these entities do not have an official one-to-one correspondence with the STIX entity types, a mapping is needed for their referencing [17]. Details on the entity types conversions can be found in Table 2.

Table 2: Conversion of ATT&CK entity types to STIX Objects.

ATT&CK Entity Type	STIX Object
Tactic	x-mitre-tactic
Technique	Attack Pattern
Mitigation	Course of Action
Group	Intrusion Set
Software	Malware/Tool

3 RELATED WORKS

The need for the automatic processing of CTI reports and retrieving entities and relations has pushed researchers to adopt Information Extraction methods in the field. However, this is not an easy task due to many reasons. First, information in raw text reports can be conveyed differently through semantics, and bulletin styles can

differ from vendor to vendor. Furthermore, reports might (and frequently do) include new entities, making the usage of a static and non-interactive database of entity templates inefficient. Moreover, ever-changing Indicators of Compromise (IOCs, e.g., IP addresses, hashes, URLs, Bitcoin addresses) must be recognized and linked to their respective actor or malware. For these reasons, many different models have been proposed to push research on constructing a Knowledge Graph in Cybersecurity [34]. One of the aspects that can be noticed in the current literature on IE techniques applied to CTI reports is that many of the proposed models focus on one or a few types of entities/relations at a time while neglecting the others. This allows researchers to obtain impressive results in a narrow domain that does not always reflect the whole needs of the CTI community. Moreover, while being very efficient, using Machine Learning and Deep Learning models in Information Extraction has a few drawbacks. First, it is not easily scalable on wider domains to include more entity types and thus requires fully retraining the model while making changes to the architecture. Secondly, the training dataset might be annotated with just a few entity types, and to include other STIX types, analysts should perform the procedure all over again. This is quite time-consuming and can become very costly for companies and organizations.

In You et al. [35], researchers have developed a model that can retrieve Tactics, Techniques, and Procedures (TTPs) with an accuracy of 0.941, which constitutes the state-of-the-art for this task. However, not only do TTPs not reflect the overall spectrum of Cybersecurity entities that should be extracted in a report, but the number of these TTPs is just 6. At the same time, the MITRE ATT&CK Knowledge Base indicates at least 14 tactics and 191 techniques (just in enterprise environments and thus neglecting mobile and ICS attacks).

Similar work has been previously done by Legoy et al. with rcATT, a Python tool used to predict MITRE ATT&CK tactics and techniques from cyber threat reports [13]. It has a maximum precision of around 0.75 in both tactics and techniques. These reduced performance levels are justified by an increased number of labels, comprising 215 MITRE ATT&CK techniques and 12 tactics. However, this work was published in 2020, and the MITRE ATT&CK framework has changed with new techniques and tactics, so careful

³<https://www.mitre.org/about/corporate-overview>

⁴<https://docs.oasis-open.org/cti/stix/v2.1/csprd01/stix-v2.1-csprd01.html>

reparametrization is needed, and (as stated in the future work’s sections of the paper) retraining it on a bigger dataset might improve its performance.

A more general approach that can tackle a broader domain of entities is the work of Gasmi et al. [9], which includes both entity extraction and relation extraction on data from the National Vulnerability Database (NVD)⁵. Researchers addressed seven common entity types and six relationship types. This database contains CVE items, i.e., disclosed cybersecurity vulnerabilities specifically formatted to be more easily cataloged, evaluated, and shared among the community. The tool reaches a precision value of 89% on the entity extraction task and 92% on the relation extraction task. However, the data used for training and testing consists of vulnerability descriptions that, while written in natural language, present a common structure and thus do not reflect the variance that might be present in more common CTI reports or bulletins.

While some proposed works leverage Machine Learning models, which perform particularly well in narrow domains, other string tagging techniques can tackle the entity extraction task.

- **Name-Matching Strings:** if a Knowledge Base containing the entity names is already in place inside the platform, it can match words inside a text. Even though the construction of the KB can be time-consuming, there are a lot of public sources from which to retrieve information on these entities, and it is also possible to provide aliases for each of them, thus being able to recognize pseudonyms and still link them to the correct entity.
- **RE-Matching Strings:** Indicators of Compromises can be found by their particular structure constant across the same type of IOC. For example, IP(v4) addresses are always be written in the format XXX.XXX.XXX.XXX, i.e., four sets of numbers from 0 to 255 separated by a dot character. Different rules apply to different types of entities. Still, in the domain of words with distinct character structures, it is possible to use regular expression tools to identify and extract them.
- **Verb-Related String:** this technique is used when the other two fail in extracting entities like companies and new malware names, which do not present any particular character structure and might not be present inside the Knowledge Base. These entities can be retrieved by analyzing each sentence through POS Tagging and Dependency Parsing. After that, we can retrieve the verb to match it with a predefined set of words that might indicate the presence of an entity.

STIXnet expands these results and further generalizes the entities and relationships that can be extracted in a document to comprise all STIX entities and relationships, which most closely represent the information an analyst should extract from CTI reports.

4 METHODOLOGY

In this section, we present STIXnet in more detail on how it works and which algorithms and models are used to perform Information Extraction on unstructured Cyber Threat Intelligence reports. We first show its pipeline and briefly overview the various modules (Section 4.1). Then we give more details on its main components: the

text extraction module (Section 4.2), the entity extraction module (Section 4.3) and the relation extraction module (Section 4.4).

4.1 Pipeline

STIXnet performs a highly complex Information Extraction task on many different types of entities and relations. There are indeed 18 different types of entities compliant with the STIX standard and more than 100 types of relations. To accomplish this, STIXnet uses different modules for each one of the different tasks that it must achieve: textual extraction, entity extraction, and relation extraction. A graphic overview of the STIXnet platform is shown in Figure 1.

- **Text Extraction:** the first module converts the program’s input into raw text. While doing this, artifacts are inevitably be introduced in the text, and therefore they must be handled to obtain a single string of characters in a common encoding.
- **Entity Extraction:** this module handles the extraction of the different entities in the text. It uses four different sub-modules to accomplish that:
 - **IOC Finder:** we extract Indicators Of Compromise using regular expression rules and text pattern extraction.
 - **Knowledge Base Entities Extraction:** using the entities in a Knowledge Base, we use an efficient algorithm for string search in a text to retrieve names and aliases. We mitigate errors and false positives caused by this approach using NLP techniques.
 - **Novel Entities Extraction:** using NLP libraries, we extract entities not present in the Knowledge Base, which can then be integrated into the KB.
 - **TTPs Extraction:** techniques and tactics are not always represented in a text by their name but can be implicit and not explicitly expressed. We use a Machine Learning model trained on MITRE ATT&CK tactics and techniques to recognize them.
- **Relation Extraction:** from the extracted entities, we now retrieve the relations. We use two sub-modules for this task:
 - **Rule-Based Approach:** using NLP techniques, we perform Dependency Parsing and compute the shortest paths between entities. Comparing the verb inside the path with the ones in the STIX relationships, we estimate the most similar one with a degree of confidence.
 - **Deep Learning Based Approach:** to adjust the results of the previous approach, we also compute embeddings from the sentences with a Deep Learning model. We then determine the similarity between these embeddings and those computed from the list of relationship labels.
- **Output:** we create a JSON file from the extracted entities and relations.

After the processing, the JSON output is rendered through a graphical interface to reproduce a graph overview of the found entities and relations. Eventual interactions with the Knowledge Base or other platform components are disclosed in the individual sections of the different modules. Indeed, we show that by interacting with the database, it is possible to improve performance over time, particularly if an analyst decides to validate the results of the STIXnet output.

⁵<https://nvd.nist.gov/>

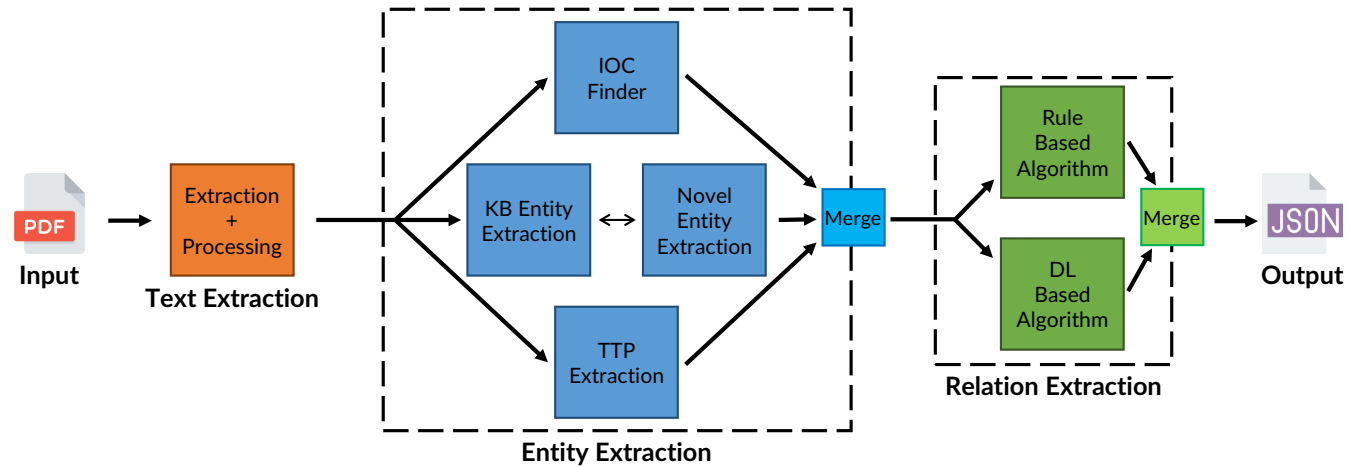


Figure 1: Pipeline of the STIXnet platform.

Furthermore, the structure of the STIXnet pipeline allows for easy and immediate management of the different modules. Formally defining the interactions between modules and submodules allows results from different pipeline components to be compared and merged in a unique output. Thus, researchers and organizations can implement the STIXnet framework in different IE scenarios and add or remove components depending on their needs.

4.2 Text Extraction

One of the platform’s most relevant and sensible aspects is its input. As mentioned, STIXnet can take in input various reports and bulletins, which can come from various vendors or sources. For this reason, reports have some stylistic and linguistic differences. However, data must be converted univocally before processing the raw text to have consistent processing between different inputs and a common ground for evaluating the various modules. This means considering many different aspects that can change from input to input. First of all, we must be able to parse text from files with different formats, and thus we use Apache Tika⁶ for pdf and doc files and ConvertAPI⁷ to extract text from the HTML data of web reports. However, since text can be formatted in many different ways, we must process it to remove artifacts, fix line breaks, and remove eventual sanitifications on IP and other addresses. This phase is crucial to ensure that the following modules are presented with a clean input; otherwise, artifacts are propagated in the pipeline and compromise the results.

4.3 Entity Extraction

This section tackles the different submodules used for entity extraction: IOC Finder (Section 4.3.1), a rule-based entity extractor (Section 4.3.2), a novel entity extractor (Section 4.3.3), and a TTPs extractor (Section 4.3.4). Finally, we clarify how the submodules interact and merge their results (Section 4.3.5).

4.3.1 IOC Finder. The particular structure of Indicators Of Compromise allows us to use regular expression (Regex) rules to find them in the text. Moreover, in some of the reports distributed by CTI vendors, the end of the document is often presented with a table containing the IOCs of interest for that particular topic. While different, all these indicators share a common structure across each type, which can be recognized with Regex rules without applying NLP techniques. Some examples of the types of IOCs supported by IOC Finder can be found in Table 3. During the execution of STIXnet and after processing the report as raw text, the first step is to run the IOC Finder submodule on it, which returns a dictionary containing the entities found.

We implement this module by forking an open-source project by Floyd Hightower⁸. To adapt it for our pipeline, we contributed to the main project by updating some libraries to a newer version and adding the ability to track the position of the found IOCs. The code of the forked project can be found in our repository.

Table 3: Examples of IOCs and their structure.

IOC Type	IOC Structure
ATT&CK Techniques	T1518 or T1518.001
CVEs	CVE-2021-44228
Email Addresses	example@mail.com
File Paths	/path/to/file
MD5s	e802c6b77dd5842906ed96ab1674c525

4.3.2 Knowledge Base Entity Extraction. Apart from IOCs, all other entity types do not share a common structure and thus cannot be found through regular expression rules. Thus, we leverage a rich Knowledge Base integrated with multiple OSINT that explicitly indicate which names represent important entities and allow us to link them to their correct entity type. For this reason, a rule-based algorithm can search specific words in the text, retrieve their

⁶<https://tika.apache.org/>

⁷<https://www.convertapi.com/>

⁸<https://github.com/fhightower/ioc-finder>

position and thus highlight them as extracted entities. The different sources for the intelligence are:

- **Knowledge Base:** Leonardo S.p.A., an Italian multinational company that collaborated in this research, provided a rich database of STIX entities. Cyber threat intelligence analysts built this database over the years at their Security Operation Center (SOC), which read and manually annotated entities from many reports.
- **MITRE ATT&CK:** the ATT&CK framework can be used as a source of intelligence for different entities such as techniques, tactics, groups, and software (of which the conversion to the STIX standard has been disclosed in Table 2). To retrieve this data, we use the Trusted Automated Exchange of Intelligence Information (TAXII) application protocol [7], which allows for exchanging threat intelligence over HTTPS and defines a RESTful API that can be used to provide or collect data.
- **Locations:** to retrieve the names of countries and continents, we used a csv file in which each nation is associated with its nationality. In this way, we can identify locations even when used as an attribute to another entity (e.g., "a Russian malware").

After retrieving the entities, a quick pre-processing is performed to unify their formats and add the possibility of aliases for each one. Through aliases, it is possible to recognize an entity in a text and map it to the correct one, avoiding duplicates and fixing the issue of multiple names for a single Advanced Persistent Threat. We then use the Aho-Corasick algorithm to find terms of this thesaurus of words in the report [1]. To mitigate false positives, we process each sentence with NLP techniques, in particular, Part-Of-Speech Tagging (POS tagging), allowing us to assign part-of-speech tags to each word (e.g., noun, verb). After defining a table of entity types and their possible POS tags, we use it to compare the entities found in the report with their extracted POS tag. For example, "us" can be used as a pronoun or can be used as a noun to reference the United States, constituting an entity of type "location".

4.3.3 Novel Entities Extraction. Some CTI reports are published to spread awareness of newly discovered actors, malwares, or techniques. These entities are thus named by CTI researchers and analysts and, for this reason, are most likely not present in the Knowledge Base. To find these new entities, we can leverage the previous execution of POS Tagging to extend its results and create a dependency graph from the tokens found in each sentence. In this way, we identify specific patterns used in the text to express a new entity. To create such a graph, we leverage both the POS tags and the dependencies between the tokens, as shown in Figure 2. To perform this processing, we use Spacy⁹, a free, open-source library for advanced NLP in Python. By looking at numerous reports and bulletins from different vendors and sources, we can identify a limited number of ways a new entity can be introduced, allowing us to write pattern rules that the NLP processing can recognize.

4.3.4 TTPs Extraction. While malwares and threat actors are often explicitly mentioned, some other entities are not and can be referenced without categorically stating their names. It is the case of tactics and techniques, which constitute the TTPs mapped into

the STIX objects "x-mitre-tactic" and "attack pattern". Rule-based methods cannot be used to retrieve these entities since the variance that characterizes the expression of these concepts is too broad and is hardly definable through a set of rules. For this reason, a multi-label text classification model must be deployed and trained on the MITRE ATT&CK Knowledge Base of tactics and techniques. We already presented a tool named rcATT [13] that suffered from its age since it was published in 2020, and the MITRE ATT&CK framework has had several changes and renovations ever since. The source code for rcATT is publicly available in their GitHub repository¹⁰, and thus we retrained it from scratch with new techniques and tactics. Also, to address one of the limitations and future works of the paper presenting rcATT, we expanded the training dataset with more data from MITRE ATT&CK descriptions and external sources for each tactic and technique. The new dataset, the pre-trained models, and the updated code for its training can be found in our repository.

4.3.5 Submodules Interaction. As stated in Section 4.1, the pipeline structure of our framework allows us to differentiate the tasks in many modules. For entity extraction, in particular, we identified four different submodules independent from one another. Their input is always the same and is constituted by the textual data of the report. Once each submodule produces an output, a final check is performed to ensure no overlapping occurs during the processing. This process includes cross-checking the found entities in the Knowledge Base to maximize their confidence in their extracted type and the addition of the novel entities. More details are disclosed in Section 5.1. Finally, all the entities are merged into one data structure, constituting the following module's input. Since this procedure is automatic, submodules can be added and removed in the pipeline according to the user's needs, and no conflicts occur during the process.

4.4 Relation Extraction

This module can retrieve relations between the found entities while processing the sentences in the raw text. However, one of Spacy's limitations is its inability to grasp relations between distant entities in a text. To address this, we propose two different approaches for relation extraction. In the first approach, we leverage POS Tagging and Dependency Parsing to compute a graph of each sentence and retrieve relations by looking at the shortest paths between entities (Section 4.4.1). In the second approach, we use a Transformer model to compute embeddings of the sentences and compute their similarity (Section 4.4.2). Finally, we clarify how the submodules interact and merge their results (Section 4.4.3).

4.4.1 Rule Based Approach. The main idea of the rule-based approach is to leverage the dependency graphs already computed by Spacy for novel entity extraction (Section 4.3.3) and use graph theory functions to grasp the relation between two entities inside a sentence. In particular, we can process the dependency graph and retrieve the relations between any couple of nodes by discovering the Shortest Dependency Path (SDP), i.e., the shortest path between two nodes in the graph. It has been observed in other studies that the nodes in the SDPs usually contain the necessary

⁹<https://spacy.io/>

¹⁰<https://github.com/vlegoy/rcATT>




Figure 2: Example of dependency graph generated from a sentence.

information to identify a relationship between two entities while also being dependent on the structure and semantic complexity of the sentence [11, 33].

After retrieving the shortest path between each couple of entities in the sentence, we extract their STIX type and focus on the verbs in the path. For example, considering the sentence in Figure 2, the extracted paths (between the three entities "APT29", "7-Zip", and "Raindrop") are:

- [APT29, used, 7-Zip];
- [APT29, used, decode, malware, Raindrop];
- [7-Zip, used, decode, malware, Raindrop].

By looking at the entity types and the root form of the verbs (underlined in the previous example), it is possible to compare them with the ones found in the list of STIX relationships and thus label the path as a specific STIX Relationship Object (SRO). However, since the verbs used to describe the relationship in the sentence might not be the same as the associated SRO, we use a similarity function to determine their likeliness. If it surpasses a certain threshold, we can consider them synonyms. To accomplish this, we use the Wu & Palmer similarity function [32], which, given two words and their synsets (i.e., groupings of similar words that express the same concept), outputs a value in the range [0, 1], where 1 means maximum similarity. For this reason, this value can be used as a confidence measure of the relation. The taxonomy used for this task is the WordNet taxonomy, an extensive lexical database of English words developed by Princeton University [8].

For example, considering the SDP [APT29, used, 7-Zip], the entity extraction module identified "APT29" as an intrusion-set and "7-Zip" as a tool. In the list of SROs, there is only one entry containing both an intrusion-set and a tool, which is "intrusion-set uses tool": since the root form of the verb in the SDP is equal to the one in the SRO, we label it accordingly with maximum confidence. Instead, considering another SDP [APT29, attacks, the, US] (where "US" is identified as a location entity), there are two SROs including both an intrusion-set and a location: "intrusion-set originates-from location" and "intrusion-set targets location". In this case, none of the root forms of SROs verbs ("originate" and "target") coincide with the one in the SDP ("attack"), and thus we compute the similarity between them:

$$wup("attack", "originate") = 0.4,$$

$$wup("attack", "target") = 0.5.$$

While keeping the confidence threshold at 0.5, we can label the SDP as a relationship of type "intrusion-set targets location".

The confidence value in the extracted relations becomes particularly important when dealing with sentences containing multiple entities. Indeed, the number of relations increases exponentially with the number of entities found, and some of the extracted relations could not exist. To include a "non-relation" label in this task, we set a threshold for the confidence, under which we discard the extracted relations.

4.4.2 Deep Learning Based Approach. The rule-based approach is particularly efficient when dealing with simple phrases or when entities are close in the graph. However, many elements might be introduced in the SDP whenever two entities are far from each other. The verb with the highest similarity could be linked to different tokens in the text and might not reach the confidence threshold. To address this problem, we use embeddings, fixed-size vectors that can also be generated from textual data by Deep Learning models such as Transformers [27].

In this specific case of relation extraction, we are interested in computing the similarity between each sentence's embeddings and the STIX relationships' embeddings. To perform these embeddings, the best tool at our disposal is Sentence BERT (SBERT)¹¹, a variation of the BERT model (Bidirectional Encoder Representations from Transformers) developed by Google AI language [26]. While BERT constitutes the state-of-the-art in many NLP applications, it becomes inefficient when dealing with a large corpus of sentence processing. SBERT addresses this problem using siamese and triplet network structures, drastically reducing processing time [22].

For its STIXnet implementation, we compute the embeddings of the different STIX relationships and the sentences extracted from the report. For each sentence, we perform a pre-processing procedure for each contained entity couple by substituting their tokens with their extracted STIX type. Then, we compute the cosine similarity between these embeddings and normalize it to use it as a confidence value. We also use a threshold for the confidence of the relation to discriminate false positives (0.5).

4.4.3 Submodules Interaction. As with the entity extraction module, the relation extraction task is divided into different submodules that work independently. While their input is always the same (i.e., textual data from the report and the previously identified entities), the two submodules perform virtually the same task this time. Thus

¹¹<https://www.sbert.net/>

we expect a degree of overlap in their results. However, their co-existence is necessary to extract relations from both simple and complex scenarios. To handle conflicts in the possible output, we use the confidence values generated during the processing and keep the relations between entities with maximum confidence, which must also be over the acceptance threshold. Therefore, users can add their desired submodules for the relation extraction task, and STIXnet will automatically merge their results.

5 EVALUATION

We now evaluate the proposed entity and relation extraction modules in STIXnet. We use three metrics to evaluate the modules: precision, recall, and F1-score. We define True Positives (TP), False Positives (FP), and False Negatives (FN) as follows.

- **True Positives:** entities or relations correctly classified by the model.
- **False Positives:** entities and relations found by the model but misclassified or that do not constitute an entity or a relation.
- **False Negatives:** entities and relations not found by the model but present in the report.

The metrics for the evaluation are Precision, Recall, and F1 Score and are defined as follows.

$$\text{Precision} = \frac{TP}{TP + FP}, \quad (1)$$

$$\text{Recall} = \frac{TP}{TP + FN}, \quad (2)$$

$$F1 = 2 \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}. \quad (3)$$

To address the lack of available annotated reports in the literature, we curated a unique dataset of CTI reports, enabling us to evaluate our model effectively. Each report treats a group or threat actor from the MITRE ATT&CK framework. All their related data has been extracted from their official ATT&CK framework descriptions and the external sources listed by the ATT&CK APIs. We then manually label each report with LabelStudio, a free and open-source software for data labeling. Both the dataset and the annotations are free to use and accessible in our repository. Annotations are exported in a JSON file and can be graphically visualized through the LabelStudio software.

To show the effectiveness of both entity and relation extraction, we tackle the evaluation of the modules separately, respectively, in Section 5.1 and Section 5.2.

5.1 Entity Extraction

While evaluating the entity extraction module for STIXnet, we must consider the different sub-modules separately since they perform different independent operations.

To have a baseline for comparing the results, we first evaluate the completeness of the deployed Knowledge Base in a fixed position in time, i.e., not being enhanced by adding new entities when found. We take the whole dataset and run the rule-based entity extraction sub-module on each report. We then extract precision, recall, and F1 scores and compute the mean of these scores over the number of

reports processed. We compare this approach with other rule-based algorithms for entity extraction found in the literature. Since rule-based approaches are domain-dependent and language-dependent, in Table 4 we compare our approach with other works in literature that operate in specific domains to more accurately resemble our task (for [4], we extracted the evaluation of the hybrid model since it more accurately resembles our rule-based algorithm). While the compared models are not designed for CTI data extraction, to the best of our knowledge, no other tools in the literature perform such an evaluation on rule-based approaches for STIX entities. As a note, the ground truth constituted by the manually annotated reports also contains novel entities and TTPs (not extractable with just this submodule), thus highlighting the contribution of the sub-module in the whole entity extraction system.

We then evaluate the novel entity extraction sub-module separately. We created a new dataset of sentences containing non-existent entities with made-up names for its evaluation. We do this since, in the original dataset, some reports do not contain any new entities and would thus bias the results of this evaluation. By creating new sentences instead, we ensure a constant number of novel entities that the sub-module can extract. In such a scenario, the isolated sub-module reaches a precision of 0.927, a recall of 0.854, and an F1 score of 0.889. Given the system’s modularity and the submodules’ specific task, a comparison with other models for Information Extraction in Cyber Threat Intelligence is given in Table 5 (the results of STIXnet extraction are obtained with the dynamically augmented Knowledge Base). Results are obtained after combining the contribution of the different approaches and evaluating them in the same dataset used for the baseline evaluation. Also, while focusing on CTI applications, we highlight the number of entity types the compared models can extract.

We also show the class-specific results for the most frequent entity types in Table 6. As we can see, the `location` class has the highest performance, given the limited number of entities and their unambiguous nature. Among the best-performing entities, there are also `intrusion set` and `malware`, thanks to their peculiar names that are easily identifiable in the text. Classes `tool` and `campaign`, however, can cause many false negatives, which is reflected by their recall values. Indeed, these concepts are fairly easy to identify once introduced in the Knowledge Base, but their novel identification might be more difficult. For example, `campaign` concepts might be hard to recognize in a text, and new tools might be referenced without their introduction (or can be confused with malwares).

5.1.1 Temporal Evolution. The interaction of the entity extraction module with the Knowledge Base allows STIXnet to easily and quickly extract previously recognized entities and constantly update the contents of the database to guarantee high-performance values over time. To ensure that the reports’ processing order does not influence the evaluation results, we randomly shuffled the dataset many times and evaluated each shuffle. We obtained a standard deviation between results close to 0 for all the evaluation metrics.

To highlight the capability of STIXnet to maintain its performance over time, we must simulate the execution of the module on subsequent reports following a temporal evolution. To do this, we evaluate its performance according to the following steps:

- (1) We divided the dataset into batches of 5 reports each.

Table 4: Baseline evaluations for the Knowledge Base entity extraction sub-module and comparisons.

Model	Domain	Entities	Precision	Recall	F1 Score
Godény [10]	Consumer electronics	Product names	N/A	N/A	0.221
Quimbaya et al. [20]	Electronic health records	Diagnosis, treatment	0.630	0.573	0.600
Chen et al. [4]	Clinical trial cohort selection	Patient data and conditions	N/A	N/A	0.845
STIXnet Rule-Based Algorithm	CTI Reports	All STIX entity types	0.835	0.869	0.846

Table 5: Comparison of evaluations for the entity extraction task in the CTI domain.

Model	Number of Entity Types	Precision	Recall	F1 Score
Weerawardhana et al. [31]	4	0.730	0.820	0.720
Li et al. [14]	4	0.839	0.789	0.813
Zhou (2022) et al. [37]	5	0.785	0.697	0.739
Zhou (2023) et al. [36]	5	0.768	0.792	0.795
Ranade et al. [21]	6	0.879	0.874	0.883
Wang et al. [29]	8	0.859	0.863	0.861
STIXnet Entity Extraction	18	0.903	0.935	0.916

Table 6: Entity extraction results for the most frequent STIX entity types.

Entity Type	Precision	Recall	F1 Score
Attack Pattern	0.702	0.861	0.771
Campaign	0.615	0.322	0.409
Identity	0.719	0.878	0.790
Intrusion Set	0.948	0.936	0.941
Location	0.962	0.913	0.936
Malware	0.888	0.790	0.835
Tool	0.949	0.560	0.698

- (2) For one of the batches, we run the entire module on each report and evaluate its performance.
- (3) Entities found by the novel entity extraction sub-modules are added to the Knowledge Base after a quick manual validation.
- (4) Repeat steps (2) and (3) for all the other batches.

We think this evaluation is fair on the premises of the possible implementations of the tool. Indeed, in a real-world application, several reports are published each day and thus need to be processed. Given the deep relationship with the Knowledge Base, we need to avoid introducing bad entities that could degrade the system’s performance in the long term. Thus, the newly extracted entities can be manually validated by an analyst, who should ensure the integrity of the data instead of fully annotating the report. The results of this evaluation are shown in Figure 3, where the evolution of the metrics is given in function of the processed batch of reports. Furthermore, to highlight the advantage of this approach concerning a static deployment of the Knowledge Base, we performed a different evaluation by repeating the same steps but skipping the third and thus freezing the state of the Knowledge Base in time. As shown, adding entities in the Knowledge Base and their quick

validation greatly improves overall performances for all three evaluation metrics and ensures the system’s updated status on the latest reports.

5.2 Relation Extraction

For evaluating the relation extraction module of STIXnet, we use the final results derived from the combined use of the rule-based approach and the deep learning based approach for relation extraction. As mentioned, relations are extracted by both submodules with a confidence value used to compare the results and eventually impose a threshold of confidence under which relations are discarded. In this way, we include the possibility of a non-relation between two entities. To compare the confidence values for both submodules, we normalized the cosine similarity of the embeddings given by the deep learning based approach in the $[0, 1]$ range. We found that a value of 0.5 for the threshold provides a fair tradeoff between false positives and false negatives.

Evaluation has been performed by comparing the relations extracted manually from the reports with the ones that the relation extraction module of STIXnet has found. This module input is constituted by the entities extracted by the entity extraction module and works on them and the sentences in the text to find possible relations. However, the modularity enforced by the STIXnet pipeline introduces the error propagation problem from entity extraction to relation extraction. Indeed, whenever an entity is misclassified by one of the sub-modules of entity extraction, it is passed as input in the relation extraction module, which inevitably produces an error since that entity constitutes a false positive. For this reason, Table 7 shows the evaluation for both scenarios. In the first one, we execute the relation extraction module after the results from the entity extraction module have been generated. In the second one, we do not consider relationships between entities where at least one was misclassified to remove the error propagation between the modules. While the F1 scores of the two scenarios are similar, the precision when removing the error propagation effect increases

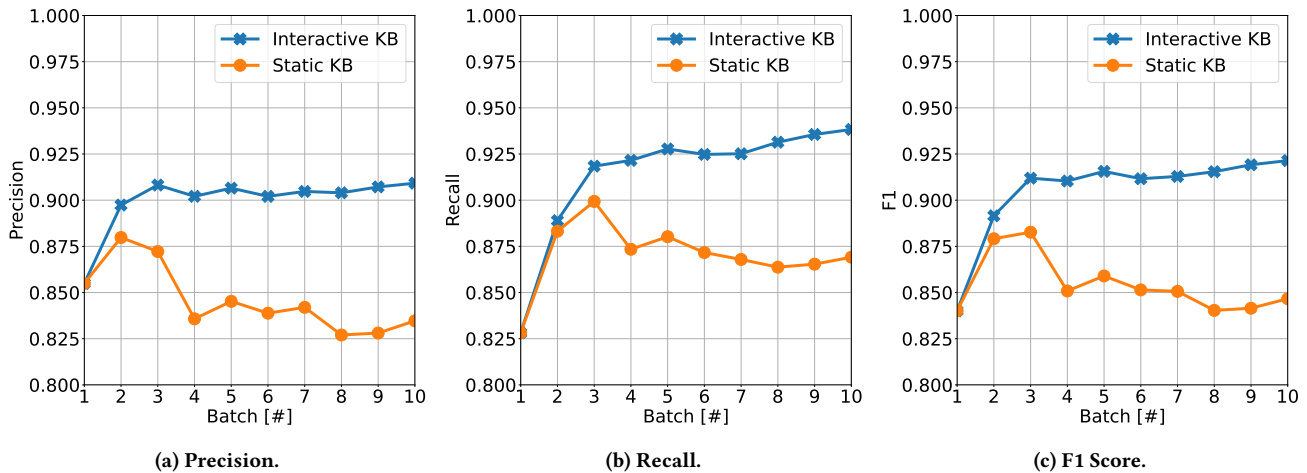


Figure 3: Evolution of the metrics for evaluating Entity Extraction in function of reports batch.

by around 0.1. This type of scenario, however, affects the recall value since the overall number of true positives is decreased, but the number of false negatives is not.

Table 7: Relation Extraction Evaluation.

Scenario	Precision	Recall	F1 Score
Standard	0.721	0.753	0.724
No Error Propagation	0.828	0.692	0.733

While the performances of the relation extraction module in the regular scenario have a lower value with respect to the ones of the entity extraction module, we must keep in mind that with each extracted entity, the overall number of possible relationships in the text exponentially increases. Indeed, to the best of our knowledge, ours is the only model that tackles both tasks subsequently by considering each of the entity types of the STIX standard and each of the STIX relationship objects.

6 CONCLUSIONS

Extracting entities and relations from CTI reports becomes more challenging as the number of classes increases. This paper presents STIXnet, the first solution for automatically extracting all STIX entities and relationships in unstructured Cyber Threat Intelligence reports. Our contribution uses rule-based, NLP, and DL techniques to retrieve all STIX entities and relationships in a report automatically. We also resort to regular expression rules and an extensible Knowledge Base to effectively create an infrastructure for cyber threat analysts to consult and gather all their data and resources. The proposed pipeline enforces a modular approach to be more flexible on the user and their demands, thus separating the different tasks into different modules. The formal definition of the interaction between the sub-modules allows researchers and organizations to use our framework by adding, swapping, and removing sub-modules at will. This is particularly useful in specific scenarios

where threat analysts might want to focus on specific entity classes, thus fine-tuning the framework to their needs.

In future works, we would like to explore the text extraction module of STIXnet further to enhance its results. Indeed, being the first module in the pipeline, accurate extraction and efficient artifact removal could improve the precision of the subsequent modules. Furthermore, STIX Domain Objects have many fields that vary depending on the entity types we are considering. All these fields can be filled with information that can be extracted from the sentences. Thus, it could be possible to extend the scope of our processing and provide additional intelligence when present.

ACKNOWLEDGMENTS

We gratefully acknowledge Leonardo S.p.A. for funding this research work and providing the contents of the Knowledge Base.

REFERENCES

- [1] Alfred V. Aho and Margaret J. Corasick. 1975. Efficient String Matching: An Aid to Bibliographic Search. *Commun. ACM* 18, 6 (jun 1975), 333–340.
- [2] Sean Barnum. 2012. Standardizing cyber threat intelligence information with the structured threat information expression (stix). *Mitre Corporation* 11 (2012), 1–22.
- [3] David Bianco. 2013. The pyramid of pain. *Enterprise Detection & Response* (2013).
- [4] Long Chen, Yu Gu, Xin Ji, Chao Lou, Zhiyong Sun, Haodan Li, Yuan Gao, and Yang Huang. 2019. Clinical trial cohort selection based on multi-level rule-based natural language processing system. *Journal of the American Medical Informatics Association* 26, 11 (07 2019), 1218–1226.
- [5] Ping Chen, Lieven Desmet, and Christophe Huygens. 2014. A Study on Advanced Persistent Threats. In *Communications and Multimedia Security*, Bart De Decker and André Zúquete (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 63–72.
- [6] K. R. Chowdhary. 2020. *Natural Language Processing*. Springer India, 603–649.
- [7] Julie Connolly, Mark Davidson, and Charles Schmidt. 2014. The trusted automated exchange of indicator information (taxii). *The MITRE Corporation* (2014), 1–20.
- [8] Christiane Fellbaum. 2010. *WordNet*. Springer Netherlands, 231–243.
- [9] Houssein Gasmı, Jannik Laval, and Abdelaziz Bouras. 2019. Information Extraction of Cybersecurity Concepts: An LSTM Approach. *Applied Sciences* 9, 19 (2019).
- [10] Balázs Godény. 2012. Rule Based Product Name Recognition and Disambiguation. In *2012 IEEE 12th International Conference on Data Mining Workshops*. 858–860.
- [11] Lei Hua and Chanqin Quan. 2016. A shortest dependency path based convolutional neural network for protein-protein relation extraction. *BioMed research international* 2016 (2016).
- [12] Natalia Konstantinova. 2014. Review of Relation Extraction Methods: What Is New Out There?. In *International Conference on Analysis of Images, Social*

- Networks and Texts*. Springer International Publishing, 15–28.
- [13] Valentine Legoy, Marco Caselli, Christin Seifert, and Andreas Peter. 2020. Automated Retrieval of ATT&CK Tactics and Techniques for Cyber Threat Reports.
- [14] Tao Li, Yuanbo Guo, and Ankang Ju. 2019. A Self-Attention-Based Approach for Named Entity Recognition in Cybersecurity. In *2019 15th International Conference on Computational Intelligence and Security (CIS)*, 147–150.
- [15] Sepideh Mesbah, Christoph Lofi, Manuel Valle Torre, Alessandro Bozzon, and Geert-Jan Houben. 2018. TSE-NER: An Iterative Approach for Long-Tail Entity Extraction in Scientific Publications. In *International Semantic Web Conference*. Springer International Publishing, Cham, 127–143.
- [16] Abhishek Nadgeri, Anson Bastos, Kuldeep Singh, Isaiiah Onando Mulang, Johannes Hoffart, Saeedeh Shekarpour, and Vijay Saraswat. 2021. KGPool: Dynamic Knowledge Graph Context Selection for Relation Extraction.
- [17] Luke Noel. 2021. RedAI: A machine learning approach to cyber threat intelligence. (2021).
- [18] Jakub Piskorski and Roman Yangarber. 2013. *Information extraction: Past, present and future*. Springer, 23–49.
- [19] Z Porkorny. 2018. What Are the Phases of The Threat Intelligence Lifecycle. *The Threat Intelligence Handbook* (2018).
- [20] Alexandra Pomares Quimbaya, Alejandro Sierra Múnera, Rafael Andrés González Rivera, Julián Camilo Daza Rodríguez, Oscar Mauricio Muñoz Velandia, Angel Alberto Garcia Peña, and Cyril Labbé. 2016. Named Entity Recognition Over Electronic Health Records Through a Combined Dictionary-based Approach. *Procedia Computer Science* 100 (2016), 55–61.
- [21] Priyanka Ranade, Aritran Piplai, Anupam Joshi, and Tim Finin. 2021. CyBERT: Contextualized Embeddings for the Cybersecurity Domain. In *2021 IEEE International Conference on Big Data (Big Data)*, 3334–3342.
- [22] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- [23] Johan Sigholm and Martin Bang. 2013. Towards Offensive Cyber Counterintelligence: Adopting a Target-Centric View on Advanced Persistent Threats. In *2013 European Intelligence and Security Informatics Conference*, 166–171.
- [24] Blake E Strom, Andy Applebaum, Doug P Miller, Kathryn C Nickels, Adam G Pennington, and Cody B Thomas. 2018. Mitre att&ck: Design and philosophy. In *Technical report*. The MITRE Corporation.
- [25] Peng Sun, Xuezheng Yang, Xiaobing Zhao, and Zhijuan Wang. 2018. An Overview of Named Entity Recognition. In *2018 International Conference on Asian Language Processing (IALP)*, 273–278.
- [26] Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. BERT Rediscovered the Classical NLP Pipeline. (2019).
- [27] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need.
- [28] Thomas D. Wagner, Khaled Mahbub, Esther Palomar, and Ali E. Abdallah. 2019. Cyber threat intelligence sharing: Survey and research directions. *Computers & Security* 87 (2019), 101589.
- [29] Xuren Wang, Runshi Liu, Jie Yang, Rong Chen, Zhiting Ling, Peian Yang, and Kai Zhang. 2022. 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, 406–413.
- [30] Rebecka Weegar. 2021. Applying natural language processing to electronic medical records for estimating healthy life expectancy. *The Lancet Regional Health – Western Pacific* 9 (01 Apr 2021).
- [31] Sachini Weerawardhana, Subhojeet Mukherjee, Indrajit Ray, and Adele Howe. 2015. Automated Extraction of Vulnerability Information for Home Computer Security. In *Foundations and Practice of Security*, Frédéric Cuppens, Joaquin Garcia-Alfaro, Nur Zincir Heywood, and Philip W. L. Fong (Eds.). Springer International Publishing, 356–366.
- [32] Zhibiao Wu and Martha Palmer. 1994. Verb Semantics and Lexical Selection. (1994).
- [33] Yan Xu, Lili Mou, Ge Li, Yunchuan Chen, Hao Peng, and Zhi Jin. 2015. Classifying Relations via Long Short Term Memory Networks along Shortest Dependency Paths. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 1785–1794.
- [34] Zhihao Yan and Jingju Liu. 2020. A Review on Application of Knowledge Graph in Cybersecurity. In *2020 International Signal Processing, Communications and Engineering Management Conference (ISPCEM)*, 240–243.
- [35] Yizhe You, Jun Jiang, Zhengwei Jiang, Peian Yang, Baoxu Liu, Huamin Feng, Xuren Wang, and Ning Li. 2022. TIM: threat context-enhanced TTP intelligence mining on unstructured threat data. *Cybersecurity* 5, 1 (01 Feb 2022), 3.
- [36] Yinghai Zhou, Yitong Ren, Ming Yi, Yanjun Xiao, Zhiyuan Tan, Nour Moustafa, and Zhihong Tian. 2023. CDTier: A Chinese Dataset of Threat Intelligence Entity Relationships. *IEEE Transactions on Sustainable Computing* (2023).
- [37] Yinghai Zhou, Yi Tang, Ming Yi, Chuanyu Xi, and Hai Lu. 2022. CTI View: APT Threat Intelligence Analysis System. *Security and Communication Networks* 2022 (03 Jan 2022), 9875199.