Graphene: Infrastructure Security Posture Analysis with AI-generated Attack Graphs

Xin Jin*†
The Ohio State University
jin.967@osu.edu

Charalampos Katsis*†
Purdue University
ckatsis@purdue.edu

Fan Sang*†
Georgia Institute of Technology
fsang@gatech.edu

Jiahao Sun†
Georgia Institute of Technology
jiahaosun@gatech.edu

Elisa Bertino
Purdue University
bertino@purdue.edu

Ramana Rao Kompella
Cisco Research
rkompell@cisco.com

Ashish Kundu
Cisco Research
ashkundu@cisco.com

ABSTRACT

The rampant occurrence of cybersecurity breaches imposes substantial limitations on the progress of network infrastructures, leading to compromised data, financial losses, potential harm to individuals, and disruptions in essential services. The current security landscape demands the urgent development of a holistic security assessment solution that encompasses vulnerability analysis and investigates the potential exploitation of these vulnerabilities as attack paths. In this paper, we propose *Graphene*, an advanced system designed to provide a detailed analysis of the security posture of computing infrastructures. Using user-provided information, such as device details and software versions, Graphene performs a comprehensive security assessment. This assessment includes identifying associated vulnerabilities and constructing potential attack graphs that adversaries can exploit. Furthermore, Graphene evaluates the exploitability of these attack paths and quantifies the overall security posture through a scoring mechanism. The system takes a holistic approach by analyzing security layers encompassing hardware, system, network, and cryptography. Furthermore, Graphene delves into the interconnections between these layers, exploring how vulnerabilities in one layer can be leveraged to exploit vulnerabilities in others. In this paper, we present the end-to-end pipeline implemented in Graphene, showcasing the systematic approach adopted for conducting this thorough security analysis.

1 INTRODUCTION

The escalating complexity of enterprise networks has resulted in environments highly susceptible to cyber-attacks [13]. The proliferation of thousands of applications deployed on network hosts and the influx of software packages from diverse sources, whether active or dormant, significantly magnify security concerns. The presence of such a diverse range of software raises questions about the potential vulnerabilities they may introduce to the network. The intricate network of interconnected devices, coupled with emerging technologies such as cloud computing, Internet of Things (IoT) devices, and interconnected systems, expands the attack surface exponentially, providing attackers with numerous entry points [26].

Moreover, the interconnected nature of networks intensifies the impact of a single vulnerability, enabling adversaries to navigate interconnected systems and compromise multiple hosts and devices. Finally, hundreds of vulnerabilities are disclosed monthly in the national vulnerability databases; thus, an approach to evaluating and understanding their impact is of paramount importance for several reasons, such as prioritizing patching efforts.

Faced with those complex challenges, ensuring a robust security posture is thus critical in today's dynamic landscape of enterprise networks, where strategies, policies, and practices collectively fortify an organization's cybersecurity defense.

Problem Scope. Given the expanding attack surface, we need comprehensive systems able not only to identify vulnerabilities specific to the infrastructure of interest but also to discern how these vulnerabilities can be exploited in sequences. Our goal is to develop a solution that leverages *attack graphs* to (1) understand how vulnerabilities might serve as a sequence of steps in a multi-step attack and (2) scrutinize and grasp the implications of a vulnerability on the infrastructure under analysis. The system implementing the solution should incorporate continuous monitoring to detect known vulnerabilities across host applications, devices, and their configurations. Additionally, a deep understanding of the nature and impact of each vulnerability is essential to tailor effective approaches for mitigation by subject matter experts.

Challenges. Designing such a system requires addressing the following challenges: **(C1)** The vulnerabilities are typically described in natural language (i.e., common vulnerabilities and exposures – CVE) rather than in a formally defined encoded format. Therefore, a systematic approach is needed to capture the vulnerability semantics and convert them into a suitable format for further analysis. **(C2)** The formulation of attack paths constituting a chain of vulnerabilities to the attacker's objectives with the least manual effort requires extracting the conditions that allow one to exploit a vulnerability (i.e., *preconditions*) and the state of the system once a vulnerability is exploited (i.e., *postconditions*). **(C3)** Finally, it is hard to come up with security quantification metrics that capture both the criticality of vulnerabilities and the impact on the system under analysis.

Our Approach. To address such challenges, we design an innovative solution, called *Graphene*, which serves as a comprehensive

---

security posture analyzer for computing infrastructures and applications. Graphene operates seamlessly by continuously monitoring trustworthy data sources, such as national vulnerability databases, for disclosed vulnerabilities specific to the devices and configuration of a given infrastructure.

Using named entity recognition (NER), Graphene is able to extract the semantic meaning of these vulnerabilities and encode them into a latent space for in-depth analysis. In particular, Graphene automatically extracts the preconditions required for an adversary to exploit a vulnerability and the result after exploiting the vulnerability (i.e., postconditions). For example, Graphene can discern that exploiting a particular CVE might require the presence of the TensorFlow XLA compiler in the Google TensorFlow version prior to 1.7.0 as a precondition. Simultaneously, Graphene captures postconditions such as a system crash resulting from a heap buffer overflow or reading from other parts of the process memory.

Subsequently, Graphene employs a semantic similarity-matching approach based on word embeddings [38] to evaluate whether the postconditions of a vulnerability match the preconditions of another, thus constructing potential attack graphs for a given topology. Once generated, Graphene analyzes these attack graphs to evaluate the security posture. We propose a novel layered approach to this analysis, where vulnerabilities within each layer share similar characteristics and can be comprehensively assessed by subject matter experts. Graphene categorizes vulnerabilities in the generated attack graphs into distinct layers, such as machine learning, system, hardware, network, and cryptography, using keyword matching and common weakness enumeration (CWE) information within the vulnerability disclosure. Hence, Graphene performs similarity matching among vulnerabilities within the same layer and across different layers, revealing how vulnerabilities in one layer can lead to the exploitation of vulnerabilities in different layers. This approach allows for prioritized risk analysis, mitigation strategies, and patching efforts based on the specific nature and severity of vulnerabilities at each layer.

As a result, Graphene automatically produces two types of attack graphs: *cumulative* (or multi-layer) attack graphs and *layered* attack graphs. Cumulative attack graphs show how an attacker could exploit Common Vulnerabilities and Exposures (CVEs) across multiple layers. In contrast, layered attack graphs focus on the exploitation of vulnerabilities within the same layer. This dual representation provides a comprehensive understanding of the potential attack paths across different layers and within individual layers.

By traversing the generated graphs, Graphene thoroughly analyzes the security posture of the entire infrastructure, providing both cumulative and layer-wise security analytics. These analytics reflect (1) the exploitability efforts required by an adversary for a step or a sequence of steps and (2) the impact of attack steps based on the criticality of the resources affected in the analyzed infrastructure. For instance, resources considered critical in a specific deployment substantially increase the impact of an attack. Once scores are computed with respect to such criticality, Graphene identifies vulnerabilities or highly exploitable and high-impact attack paths that require immediate attention to safeguard the system from potential attacks.

We make the following contributions:

- We introduce a novel fully-automated security posture analyzer designed to generate attack graphs for computing infrastructures.
- We propose a natural language processing approach based on NER and word embeddings that streamlines pre- and postcondition extraction, facilitating the generation of attack graphs without the need for manual intervention.
- Our framework adopts a comprehensive strategy for analyzing security postures in a multi-layered fashion. The approach analyzes each layer separately and combines them into one unified analysis.
- We propose risk scoring methods for tailored analysis of the underlying network infrastructure.

Roadmap. Section 2 motivates Graphene. Section 3 provides background information on vulnerability disclosures and on natural language processing using NER. Section 4 gives an overview of the Graphene pipeline, while in Section 5, we delve into the technical details, starting from the ML-based processing to the attack graph generation. Subsequently, Section 6 outlines the risk analysis methods incorporated in Graphene, designed to evaluate the security posture of a given infrastructure. In Section 7, we provide a comprehensive discussion of the implementation details, followed by the presentation of the evaluation results in Section 8. Finally, Section 9 discusses the related work and Section 10 concludes the paper.

2 MOTIVATION

In this section, we delve into the rationale behind employing attack graphs as a fundamental tool for conducting posture analysis. We then discuss limitations of previous approaches - limitations that are addressed by Graphene.

Security Posture Analysis Using Attack Graphs. To achieve a holistic security assessment in complex systems, attack graphs can serve as a valuable tool. They provide a comprehensive representation of the system's security landscape, accounting for the intricate interdependencies among its diverse components and potential pathways for attacks. On the one hand, attack graphs give a holistic overview of the prevailing threats confronting the system, providing security professionals with a streamlined understanding of the overall security posture. On the other hand, attack graphs can facilitate the creation of simulation environments, enabling the exploration of hypothetical scenarios when addressing threats. Such a simulation environment essentially functions as a "threat lab," empowering security professionals to optimize resource allocation for mitigating threats, thus making the threat intelligence processes more cost-effective.

Additionally, security professionals can simulate diverse scenarios by manipulating the attack graphs. For instance, the deliberate disconnection of paths through the removal of a node or edge offers insights into the potential value of investing resources to address specific threats. Conversely, the addition of a node or edge triggers the automatic inference of corresponding graph components, thereby showing the repercussions of novel threats. This proactive approach aids in better preparation for future challenges and the formulation of response defensive actions and protocols. In instances where security professionals possess insights beyond

publicly available information, they can update the details of a node or edge, infusing human experience into the threat analysis process and fostering a more precise representation of the actual system.

Limitations of Prior Works. While prior research has explored various approaches to generate attack graphs, their practical application is limited by several factors. Some approaches require providing the relevant infrastructure-related Common Vulnerabilities and Exposures (CVEs), which in turn may require manual input [4, 6, 23]. However, acquiring information on infrastructure vulnerabilities must be an ongoing process, as new vulnerabilities are continually disclosed in public databases. A single vulnerability can substantially alter the security landscape by introducing new attack possibilities or high-risk pathways.

Automation is another significant limitation. Some approaches mandate that vulnerabilities be defined using proprietary formal expressions [4, 6, 15]. These formal definitions may be highly specific to a particular domain and must be regularly updated to accommodate emerging attack vectors. The quality of the resulting attack graph is intrinsically linked to the accuracy of these definitions.

Moreover, a critical issue in attack graph generation is the reliance on hard-coded heuristics and keyword matching to identify pre- and postconditions, which are essential for connecting CVEs within the attack graph [4, 15, 23, 58]. Those approaches may not be widely applicable to all CVEs and require continuous updates. Additionally, they cannot capture semantic information embedded in natural language text, which is essential for reasoning about the interconnection of CVEs.

Finally, many of the suggested risk analysis methods neglect to assess risk with respect to the underlying infrastructure [3, 24, 41, 42]. This represents a notable limitation, given that the consequences of an adversary exploiting a vulnerability in a critical resource, such as a database containing sensitive client information, are considerably more severe than those associated with a less significant resource. Consequently, conducting a vulnerability assessment tailored to the specifics of the infrastructure is imperative for fully understanding and addressing its potential impact.

3 BACKGROUND

3.1 CVE and CWE

Common Vulnerabilities and Exposures (CVE) disclosures form a critical aspect of the cybersecurity landscape, offering a standardized approach to identifying and cataloging known vulnerabilities. Assigned unique identifiers, CVEs enable effective communication and collaboration within the cybersecurity community, facilitating the timely detection and mitigation of potential threats. Concurrently, the Common Weakness Enumeration (CWE) system complements CVE by classifying and categorizing the underlying weaknesses that lead to vulnerabilities, providing a structured framework for understanding and addressing security issues. Adding to this ecosystem is the Common Vulnerability Scoring System (CVSS), which assigns numerical scores (ranging from 0 to 10) to CVE entries, quantifying the severity of vulnerabilities based on factors such as impact on confidentiality, integrity, and availability.
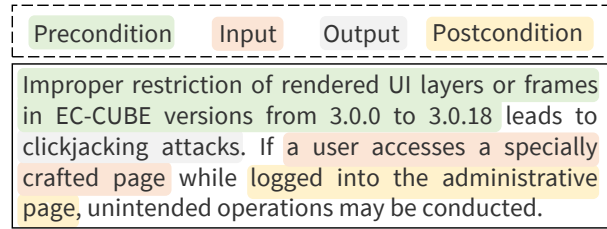

Figure 1: An example of Vulnerability Description and Attack Graph Node Attributes for CVE-2020-5679.

3.2 Attack Graphs

At the high level, an attack graph can be defined as a structured representation of the potential paths an attacker can take to compromise a network or system by exploiting vulnerabilities [4, 24]. In such a representation, the *attack graph nodes* are the attack units that are the assembly of basic vulnerability attributes combined as coherent entities. These attributes encompass preconditions, postconditions, inputs, and outputs. We give the notion of each attribute as follows.

- **Precondition**. Preconditions refer to a collection of system properties that must hold for an exploit to succeed. If these preconditions are met, it becomes possible to render all subsequent steps of an attack ineffective.

- **Postcondition**. Postconditions are the system properties that hold as the results of an attack step, which are necessary for the generation of outputs.

- **Input**. The inputs are the actions that attackers need to take to trigger the vulnerability and perform the exploit. For instance, a victim software declares a fixed-length buffer to exploit a buffer overflow vulnerability. The input for such a vulnerability is access to the buffer elements that exceed the buffer length. The vulnerable system with the attack precondition can still execute normally without attack inputs. For example, the victim program with a boundary check to prevent buffer overflow is still safe if the adversary does not perform the out-boundary element operations.

- **Output**. The outputs refer to the final values or results that the system returns or produces when exploits to vulnerabilities are executed. For example, the output for the buffer overflow vulnerability can be a system crash, abnormal privilege escalation, remote code execution, etc.

Attack graph nodes can be categorized as attackers, attack targets, and vulnerabilities. Generally, attackers act as the source or root nodes, vulnerabilities function as intermediate nodes, and attack targets represent the sink/leaf nodes.

By exploiting vulnerabilities, attackers can perform a sequence of steps or actions to attack the victims, e.g., gaining unauthorized access to a network or system. Such steps and actions are the edges in the attack graph, serving as the basic connections. Formally, *attack graph edges* represent the transitions and chains of vulnerabilities [48]. Successful attacks often require the execution of a series of exploits in a specific sequence. For example, in order to compromise the macOS Kernel through Safari, the exploit of a chain

of six vulnerabilities is required [30]. Therefore, the edge connecting two nodes in an attack graph indicates that the vulnerability exploited by one node can serve as the input and trigger for the other node.

3.3 Named Entity Recognition

Named entity recognition (NER) involves the identification of various segments of information within a text and subsequently categorizing them into predefined classes [46]. These classes can encompass a wide range of entities, including but not limited to individuals, organizations, and geographical regions [53]. In its early stages, NER heavily relied on deterministic rules [18]. However, this rule-based approach often fell short of the desired efficiency standards [35]. One significant breakthrough came with machine learning algorithms, which introduced probabilistic approaches to entity recognition [22, 35, 65]. For instance, popular learning-based NER techniques include the hidden Markov model [66] and conditional random field [53]. More recently, deep learning methodologies have played a pivotal role in enhancing the performance of NER systems. Notably, Xu et al. introduced a novel approach by combining the bidirectional LSTM model with CRF [65].

3.4 Word Embeddings

Word embedding is a natural language processing technique that represents words in a high-dimensional vector space, where each dimension corresponds to a specific feature of the word [43, 44]. The vector values are learned through a mathematical process that analyzes the co-occurrence patterns of words in a large corpus of text data so that words that occur in similar contexts are mapped to vectors that are closer together in the vector space, while words that are dissimilar in context are mapped to vectors that are farther apart. Many pre-trained word vectors, such as word2vec and Glove [47, 50], have been used in a variety of natural language processing tasks, such as text classification, machine translation, and named entity recognition.

4 SYSTEM OVERVIEW

In this section, we present an overview of the Graphene pipeline, followed by a discussion of the system's architecture.

Figure 2 provides a comprehensive overview of the Graphene pipeline, encompassing the following key steps:
(1) Data Curation: This initial stage involves utilizing user-provided details about the network infrastructure, including information on network topology, communicating entities, and device specifications. The system then queries the national vulnerability database to retrieve Common Vulnerabilities and Exposures (CVE) disclosures that are specific to the installed components, such as applications and devices.
(2) ML Processing: In this phase, the system leverages machine learning (ML) techniques to extract contextual vulnerability information from natural language text automatically. This includes capturing pre- and postconditions. The system employs word embedding techniques to establish vulnerability similarity based on the extracted conditions. Additionally, it classifies each vulnerability node to its corresponding layer based on keyword matching and CWE information provided within the CVE disclosure.
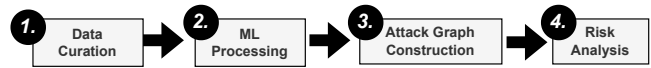

Figure 2: Graphene' pipeline.

(3) Attack Graph Construction: In this phase, Graphene constructs attack graphs by integrating the attack graph nodes identified in step 1 with the vulnerability similarity results obtained in step 2.
(4) Risk Analysis: Finally, Graphene conducts a thorough security posture analysis based on the attack graphs generated by the previous step. Graphene traverses the graphs, utilizing various vulnerability parameters such as exploitability and risk scores. This quantification process is essential for assessing the security posture layer by layer and cumulatively. Additionally, it helps identify the most impactful attack paths within the network infrastructure.

5 AUTOMATED ATTACK GRAPH GENERATION

In this section, we present our approach to building attack graphs based on machine learning and natural language processing. The first step of the approach is the identification of the fundamental components of attack graphs, i.e., graph nodes, by using named entity recognition – discussed in Section 5.1. The second step is to establish edges between attack graph nodes by modeling preconditions, postconditions, inputs, and outputs using word embeddings – discussed in Section 5.2.

5.1 Attack Graph Node Identification

The attack graph nodes are the fundamental element of attack graphs, where they provide relevant vulnerability details, i.e., preconditions, postconditions, inputs, and outputs. In contemporary computer systems, numerous attack surfaces exist, amplifying the scope of extensive attack graph nodes. As highlighted in Section 2, current approaches have significant shortcomings in scalability and adaptability. For instance, a previously proposed approach to identify attack graph nodes [48] relies on network sniffing to detect these nodes, which only applies to IoT networks. We address such limitations by analyzing vulnerability descriptions and reports, thus enabling large-scale comprehensive analyses.

Identifying attack graph nodes from vulnerability descriptions and reports is not trivial, as it requires understanding the semantics of vulnerabilities. Consider the description of CVE-2020-5679 shown in Figure 1, which can be exploited to perform clickjacking attacks. With manual efforts, human analysts can identify the key properties of this vulnerability by understanding the description and identifying the words or phrases that represent the nodes. For example, the precondition of CVE-2020-5679 consists of the target version of EC-CUBE with the improper restriction of UI layers. To identify such a precondition, we need to interpret the semantics of this description. A naive approach is to manually identify it. However, manual identification of graph nodes based on human knowledge is not scalable. Another approach is to use pivot words [58] or heuristic rules [4, 10, 16, 23, 48]. For example, the word "by" is used to identify vulnerability conditions as pivot
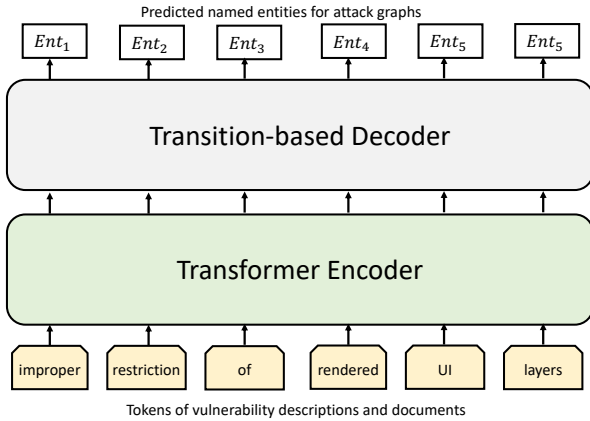
Figure 3: The Named Entity Recognition Model.

Figure 4: The Security Corpus Curation Framework.

words [58]. However, these methods are not generalizable since natural language is noisy [29], e.g., different vulnerability descriptions can include semantic but syntactically different expressions.

Node Identification by Entity Recognition. We propose to use NER for automatic attack graph node identification, which is a process of identifying entities in input texts by classifying words and phrases in vulnerability texts into the corresponding entities (addressing **C1**). While existing approaches have applied NER to analyze security reports [8, 28, 62], they either only focus on specific domains, e.g., home computers [8], or use weak models [28] (see Section 9 for details). Observing such limitations, we propose identifying security entities with a more general scope and using generative language models. In Figure 3, our model takes as input the vulnerability descriptions and documents, which are tokenized into natural language tokens. The transformer encoder is a pre-trained model (e.g., roBERTa model [40]), which has been pre-trained on a very large corpus. The encoder generates semantic embeddings by mapping the input tokens into latent space. The transition-based decoder maps the token embeddings into named entities based on a finite-state transducer [34].

In addition, NER involves the other crucial task of identifying the entity type and semantics. However, existing solutions focus on entities like software versions and network ports [8], which are often limited to particular domains and lack general applicability. Meanwhile, the MITRE CVE project released the official template for *all CVE descriptions* that covers the critical details for automated phrasing [31]. Therefore, we define six entities based on this template: vulnerability type, affected product, root cause, impact, attacker type, and attack vector, which are general to vulnerabilities because of the wide adoption of the MITRE CVE template. After identifying the various entities, we systematically categorize them as attributes of attack nodes, namely preconditions, postconditions, inputs, and outputs. In this framework, we classify the affected product entity as a precondition. The vulnerability type is categorized as a postcondition. Attacker type and root cause are considered inputs. Finally, the impact and attack vector entities are categorized as outputs. In cases where certain vulnerabilities, such as public CVEs, have been assigned manual evaluation scores such as exploitability, severity, and impact score, we can also incorporate these scores as a "golden standard" for the corresponding attack graph nodes, if available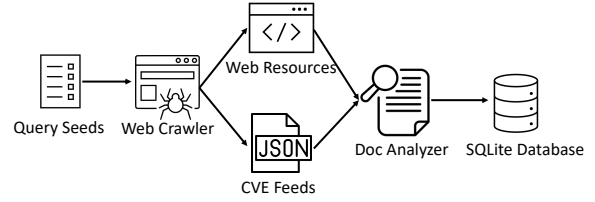. This additional information enhances the precision and accuracy of the attack graph by incorporating established metrics for attack graph analysis.

5.2 Attack Graph Edge Connection

The edge connection between two graph nodes indicates that the vulnerability exploited by one node can serve as the input and trigger for the other node. To establish connections between attack nodes, it is necessary to identify pairs of nodes whose input and output ports can be matched under the same preconditions and postconditions. Since such node attributes are identified from natural language descriptions, a simple way to build edges is to match the words of different graph nodes. However, such a solution may not always work since morphological words (e.g., synonyms, abbreviations, and misspellings) are widely used in vulnerability descriptions and texts. For instance, in the case of web injection vulnerabilities, the abbreviation "XSS" is frequently employed to refer to cross-site scripting. To overcome such limitations, we propose to utilize word embeddings to semantically match attack graph nodes and construct attack edges to address the requirement of precise node matching.

Word Embedding-based Edge Construction. In natural languages, words in different contexts can carry different meanings. For example, the word "band" has different meanings under material and music context and corpus. In other words, the word embeddings trained on one specific domain cannot be directly used in other domains because there can be semantic changes. Therefore, we argue that existing word vectors (e.g., Glove and word2vec [47, 50]), which are pre-trained on general English corpus (e.g., Wikipedia), cannot precisely deliver word semantics in the security domain.

In order to obtain embeddings tailored for holistic security assessment, we propose to train embedding models using a corpus specifically focused on security (addressing **C2**). This approach offers two distinct advantages: (1) The resulting word embeddings will facilitate the accurate semantic matching of attack nodes, allowing for precise identification and classification. (2) By quantifying the matching outcomes as similarity scores, we can assign weights to the attack edges, enabling a more nuanced representation of the severity or relevance of each attack. To obtain the similarity scores, we first compute the semantic representations of attack node attributes (e.g., preconditions, postconditions, inputs, and outputs) by averaging the vectors of all words in the ports. Next, we calculate the node similarity by the cosine similarity of the attribute semantic representations.

The key task for the word embeddings for GRAPHENE is to curate the security corpus that is general to different vulnerabilities and
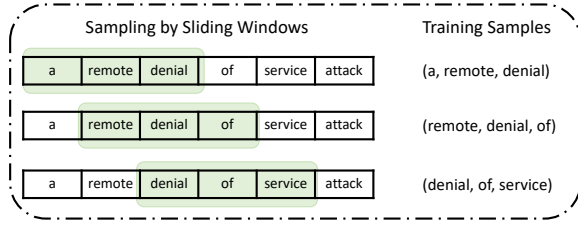
Figure 5: The Dataset Sampling and Training Samples.

attack scenarios. As discussed in Section 3, the NVD database contains detailed vulnerability descriptions and references of CVEs that have undergone thorough manual assessment and processing [9]. This has resulted in the generation of succinct information regarding various vulnerabilities. For the CVE description collection, we download CVE feeds hosted on the NVD website, which package CVE entries with descriptions. Moreover, there are also many other online resources, such as MITRE ATT&CK knowledge base [56], that collect comprehensive vulnerability information such as Common Vulnerability and Exposures (CWEs), which we also consider in our dataset curation process.

We build a security corpus curation framework as shown in Figure 4. The framework takes as input a list of query seeds (e.g., CVEs and keywords). The output of the web crawler includes online web pages and resources and the CVE feeds. Note that the CVE feeds contain much useful information, such as CWEs, vulnerable products, affected versions, CVSS scores, and reference links [14]. After obtaining all the online resources and CVE feeds, we run the doc parser to clean up these documents and extract useful content, such as the CVE descriptions and vulnerability texts.

To clean our curated dataset, we first process the dataset by removing punctuation. We do not follow the common natural language processing practice of removing stop words, as some of these words play a key role in the security context. For example, removing the stop word "of" from the phrase "denial of service" will break this popular attack phrase and change its semantics. Next, we sample the processed security documents into training samples using sliding windows as shown in Figure 5. The training samples are sets of center words and context words. For example, in the word set (a, remote, denial), remove is the center word while a and denial are the context words. This sampling method enables us to learn the word semantics by modeling the context information. For example, the word embedding model Continuous Bag-of-Words (CBOW) [43] learns word semantics by optimizing the probability likelihood estimation:

$$\frac{1}{N} \sum_{i=1}^{N} \sum_{-c \leq j \leq c, j \neq 0} log(p(w_i \mid w_{i+j})) \qquad (1)$$

where N is the total number of words in the security document, c is the sliding window size, and w_i and w_{i+j} are the center word and context word.

5.3 Attack Graph Construction and Partition

The attack graph comprises three primary node types: the attacker (source nodes), existing CVEs associated with different network

Table 1: The "Keywords" column corresponds to high-frequency keywords found in network vulnerability listings. The "Protocol" column corresponds to communication protocols between entities.

Keywords	Protocols
access control	tls
authentication	ssl
authenticity	tcp
authorization	ip
availability	http
botnet	https
CDN	ftp
certificate	ftps
certificates	udp
client	lte
cloud	wifi
communication protocol	

entities (intermediate nodes), and CWEs serving as attacking targets (sink nodes). The construction of edges connecting nodes and the assignment of corresponding edge weights depend on the types of connected nodes and available data. By default, an edge is established from the attacker node to each CVE node, assuming the attacker can exploit the CVE. This inclusivity aligns with Graphene' objective to generate all conceivable scenarios under various adversarial assumptions in the network. For example, if a CVE's attack vector necessitates physical access to the device, Graphene includes an analysis of such a scenario. Users can subsequently filter the generated attack graphs based on their network's specific adversarial assumptions. The edge weights for these edges are predominantly determined by CVSS base scores, indicating the likelihood of a successful attacker exploit, directly correlating with the severity of the threat posed by exposure to such vulnerabilities.

For any given pair of CVE nodes, an edge is established if the postcondition of one CVE node aligns (partially) with the precondition of another CVE node, indicating the potential for an attacker to exploit one vulnerability to access the other—essentially forming a chain of vulnerabilities for exploitation. Beyond relying on CVSS scores, the weights assigned to these edges are contingent on the node matching score of the two nodes, derived from the word-embedding-based node matching. The node matching score gauges the semantic similarity between the postcondition and precondition of the directed edge connecting two CVE nodes, offering insights into the difficulty an attacker might face in reaching the latter CVE by exploiting the former. Additionally, to mitigate graph complexity arising from marginally related CVE nodes, the user can stipulate the construction of edges only when the node matching score exceeds a specified threshold. Consequently, the resulting attack graph comprises paths that are more feasible for attackers, optimizing computational resources by excluding less viable paths. Lastly, each CVE node is linked to CWE nodes, representing the system's encountered threats. In the CVSS database, each CVE is associated with one or several corresponding CWEs, each instantiated as an edge in the attack graph. Similar to other edge types, the CVSS scores determine the edge weights, with the option to prune edges by setting a weight threshold.

Table 2: The table shows a list of the related CWEs pertaining to network security vulnerabilities. The "CWE ID" is a unique weakness identifier, and "CWE Name" provides more information on the identifier.

CWE ID	CWE Name
20	Improper Input Validation
79	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting')
80	Improper Neutralization of Script-Related HTML Tags in a Web Page (Basic XSS)
83	Improper Neutralization of Script in Attributes in a Web Page
87	Improper Neutralization of Alternate XSS Syntax
89	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection')
90	Improper Neutralization of Special Elements used in an LDAP Query ('LDAP Injection')
91	XML Injection (aka Blind XPath Injection)

Attack Graph Partition. After obtaining the holistic security posture, another potential need is to get an in-depth security analysis by focusing on the specific layer or scope of the target systems. To achieve this, we propose the attack graph partition component to get subgraphs from the cumulative attack graph.

Graphene follows two approaches to classify the identified vulnerabilities to the layers of interest. In the first approach, Graphene uses a predefined set of high-frequency keywords that typically appear in the corresponding layer. Table 1 shows a snippet of the keywords Graphene uses for mapping the vulnerabilities to the network layer. We list the complete list of keywords in Appendix A. The second approach is based on the MITRE Common Weakness Enumeration (CWE) identifier appointed by the vulnerability database. Table 2 shows a snippet of the list of CWEs Graphene uses for mapping the vulnerabilities to the network layer. We list the complete list in Appendix A. The reason for using both approaches is that, on the one hand, only keyword matching may result in misclassified vulnerabilities. On the other hand, relying merely on the CWEs is not enough, as many vulnerabilities have not been appointed to CWEs. Additionally, the Dashboard service allows an administrator to change the layer to which a particular vulnerability has been assigned.

6 RISK SCORING SYSTEM

Following the construction of attack graphs, a comprehensive assessment of the security posture is conducted through a risk-scoring system (addressing **C3**). This system receives the attack graphs generated by Graph MS and produces analytics on the security posture.

Within the graphs, each node (i.e., CVE) is linked to exploitability, impact, and risk scores. These scores are derived in accordance with the Common Vulnerability Scoring System (CVSS) standard [17] (version 3.1). Recognizing that these scores offer a limited, isolated perspective on vulnerability impact, Graphene enhances the computed scores based on CVSS standards and integrates them for the risk assessment of attack graphs, considering the criticality of affected resources. The evaluation of security posture involves:

(1) Computation of exploitability, risk, and impact scores for each graph.
(2) Identification of the shortest paths.
(3) Identification of the high-severity paths.
(4) Identification of the key vulnerabilities requiring immediate patching.
(5) Determination of the minimum set of vulnerabilities covering all attack paths.

In what follows, we describe how this risk-scoring system achieves the aforementioned points.

Computing graph exploitability, risk, and impact scores. The first step is to compute the scores based on the attack graphs.

Edge Exploitability Score (EES). Let ees_i denote the EES of the edge i. It is computed as follows:

$$ees_i = eScore(source(i)) + c \sum_{x \in in\_edges(source(i))} ees_x \quad (2)$$

In the above expression, $source(i)$ is a function that returns the source node of edge i while $eScore()$ returns the exploitability score as a function of the node provided as an argument. The function $in\_edges()$ returns all inbound edges to the node provided as an argument, and c is a predefined multiplication constant ($c = 0.1$ in the experiments). Hence, the calculated EES for each edge is associated with the exploitability score of the edge's source node and the scores of all preceding edges in all the paths that include the edge. The computation of ees_i requires all the EES of all edges in the set $in\_edges(source(i))$ to be calculated first.

Edge Impact Score (EIS). Let eis_i denote the EIS of edge i. It is computed as follows:

$$eis_i = iScore(sink(i)) + k \cdot \sum_{x \in out\_edges(sink(i))} eis_x \quad (3)$$

In the above expression, $sink(i)$ is a function that returns the sink node of edge i while $iScore()$ returns the impact score as a function of the node provided as an argument. The function $out\_edges()$ returns all the outbound edges from the node provided as an argument, and k ($k = 0.01$ in the experiments) is a predefined multiplication constant. Hence, the computed EIS of each edge is associated with the impact score of the edge's sink node and the scores of all the subsequent edges in all the paths that include the edge. In simpler terms, the EIS score of an edge accumulates the impact of all subgraphs starting with that edge. The computation of eis_i requires all the EIS of all the edges in the set $out\_edges(sink(i))$ to be computed first.

The functions $eScore()$ and $iScore()$ within Graphene are intentionally designed to be user-provided. Formulating a precise function that accurately encapsulates the exploitability and impact scores across diverse deployment scenarios poses a considerable challenge. For instance, the impact of a CVE affecting a device deployed within a critical infrastructure network significantly differs from its impact when the same device is located in a publicly-facing demilitarized zone. To address this complexity, users are granted the flexibility to define their own customized functions. These functions can incorporate the assigned CVE scores along with other parameters of their choosing, allowing for a more tailored and context-aware evaluation of exploitability and impact in varied deployment scenarios. If those functions are not provided, Graphene uses each CVE's assigned CVSS exploitability and impact scores.

Edge Risk Score (ERS). Let ers_i denote the ERS of edge i. It is computed as follows:

$$ers_i = weight_i \cdot (ees_i + eis_i) \quad (4)$$

In the above expression, $weight_i$ is the weight of edge i, and the values ees_i and eis_i are computed as per the Equations 2 and 3, respectively. Hence, the computed ERS of each edge is associated with

the impact and exploitability scores of the edge i. The computation of ers_i requires all the EIS and EES of the edge i to be computed first.

Once the computation of the edge scores is completed, we normalize them on a scale of 0 (low) to 10 (high). We apply the normalization for each set of edge scores (i.e., EES, EIS, ERS) for each score as follows.

$$normalized\_edge\_score = \frac{10 \cdot edge\_score}{max\_set\_score} \quad (5)$$

In the above equation, $normalized\_edge\_score$ is the calculated normalized score for a particular edge, $edge\_score$ is the original edge score, and max_set_score is the maximum score in the set. *Graph Scores.* For each generated graph, GRAPHENE computes the graph exploit, impact, and risk scores. Each score is computed as the average of the EES, EIS, and ERS sets.

Identifying the shortest paths with respect to the attacker goals. We define the shortest path towards the attacker's goals as the path in which the sum of the edge scores present in the path is the highest. Given a graph, we first find the maximum exploitability score of all the nodes in the graph. We, then, define the weight of the edge i as follows:

$$edge\_weight_i = max\_exploit - eScore(source(i)) \quad (6)$$

In the above equation, $max\_exploit$ is the maximum exploitability score of all the nodes in the graph. Thus, the higher the exploitability score of the source node of an edge i, the lower the weight assigned to the edge. We then add a temporary node to the graph where all the sink nodes (i.e., the attacker goals) have an edge targeting it. This node is called the "supersink" node. Finally, we run a weighted shortest path algorithm from the attacker node to the sink node. The algorithm runs in polynomial time and finds the paths with the least weight, thus the highly exploitable paths.

Identifying the high severity attack paths. To compute all the paths from the attacker node to the attacker's goals, GRAPHENE searches for all the possible paths up to a certain number of edges; this is the cutoff limit for path explorations. For each computed path, GRAPHENE computes the exploitability, impact, and risk score of every path, which is essentially the total sum of EES, EIS, and ERS of the edges in the path. Then GRAPHENE sorts the paths in descending order of risk, exploitability, and impact. GRAPHENE allows the system administrator to change how sorting is performed through the dashboard interface.

Identifying the key vulnerabilities that require immediate patching. To find those vulnerabilities in the attack graph, we measure the degree of every intermediate node (i.e., any node other than the source or sinks) in the graph. The degree of the node is defined as the number of edges that are incident to the node.

A node with a high degree essentially means that the node (i.e., CVEs in our attack graphs) is present in several attack paths. Thus, eliminating (i.e., patching) such nodes may render several attacks impractical.

Identifying the minimum set of vulnerabilities that cover all the attack paths. We apply a minimum set of vertex cover on the constructed attack graph to identify the minimum set of vulnerabilities that cover all the edges in the graph. As vertex cover is an NP-hard problem, we use a local-ratio approximation algorithm to find the minimum set vertex cover [5]. Thus, GRAPHENE
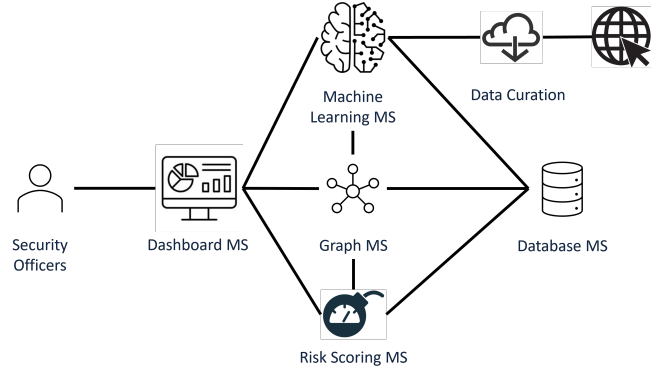
Figure 6: GRAPHENE is composed of five microservices (MS): the Dashboard MS, Graph MS, Machine Learning MS, Risk Scoring MS, and Database MS.

identifies the minimum set of vulnerabilities that could render the attack paths impractical once mitigated.

7 IMPLEMENTATION

We have implemented GRAPHENE of 6,116 lines of Python codes. GRAPHENE is comprised of five microservices (MS) (see Figure 6). Steps 1 and 2 of the pipeline (see Figure 2) are implemented by the Machine Learning MS, step 3 by Graph MS and step 4 by the Risk Scoring MS. Those services are implemented with 987 lines of Python code. Specifically, we build the named entity recognition and word embedding models based on spaCy [21], Gensim [51], and Scikit-learn [49]. For the risk scoring system, we use the NetworkX [20] library for graph construction, graph traversal, and risk score computation.

The Dashboard MS orchestrates the entire GRAPHENE pipeline. It has been developed with about 1.5K lines of code using Flask [2] and Dash [1] libraries. The dashboard provides an interface allowing the user (e.g., security officers and network admins) to import details about the network infrastructure, such as the network topology, communicating entities, and device details. Once the posture analysis is completed, the dashboard presents the generated graphs and the results of the risk analysis. For instance, Figure 8 shows GRAPHENE' front-end dashboard interface snapshot. The figure showcases a spider graph delineating the top three identified vulnerabilities and subsequently outlines the top three high-severity paths. Additionally, a table excerpt is presented, detailing identified vulnerabilities along with their corresponding extracted pre- and postconditions.

The user can also introduce new nodes in the graph (i.e., new CVEs) and modify or delete existing ones. Such functionality is needed as often online resources are not immediately updated once a certain vulnerability is disclosed. Moreover, the risk, exploitability, and impact scores may not reflect the criticality of the resources in the infrastructure (e.g., overestimated or underestimated scores). Thus, GRAPHENE allows users to change those scores. Once the user makes such changes, the dashboard communicates the request to the back end, where a change-impact analysis occurs. The results are returned to the dashboard, which is immediately updated.
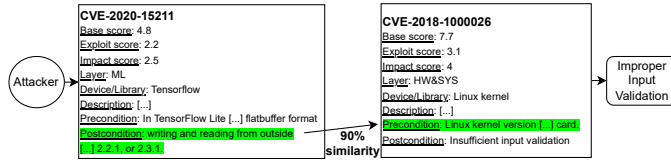
Figure 7: An example of a generated attack graph with two intermediate nodes; the first node is an ML layer CVE, and the second is a hardware & system CVE.

The dashboard also illustrates the generated attack graphs for the given infrastructure for the different layers. For example, Figure 7 shows an example of a multi-layer attack graph. The first CVE node is classified under the ML layer, whereas the second CVE under the hardware and systems area.

The dashboard also provides visual representations of the generated attack graphs pertaining to distinct layers within the designated infrastructure. For instance, Figure 7 shows a multi-layer attack graph generated by GRAPHENE. In the example in the figure, the initial CVE node is categorized within the Machine Learning (ML) layer, while the subsequent CVE node is categorized within the domain of hardware and systems.

The Databases MS acts as a store and query module. All the constructed graphs and their analysis are stored in the Database MS. This MS uses a high-performance graph database management system (i.e., Neo4j graph database [45]) to expedite querying, storing, and updating operations for large and complex attack graphs.

8 EVALUATION

In this section, we first evaluate the effectiveness of our NER-based model for processing CVE disclosures, followed by an evaluation of our word embedding technique used for attack graph construction. We then present a case study reasoning about the output of GRAPHENE. Our evaluations aim to answer the following research questions:

- **RQ1**: How effective is the named entity recognition model for attack graph nodes?
- **RQ2**: Can the proposed word embedding method semantically match the attack graph nodes?

Test Environment. We deploy GRAPHENE on a Google Cloud virtual machine with an Intel Broadwell CPU, 4 GB memory, 110 GB storage, the Ubuntu 20.04 operating system, and an NVIDIA Tesla P100 graphics card.

8.1 RQ1: Effectiveness of Named Entity Recognition Model

To answer **RQ1**, we evaluate our named entity recognition (NER) model by comparing it with baseline models on a large vulnerability NER dataset.

Dataset. Our dataset is adapted from PMA [19], which contains descriptions of 52,532 CVEs, 245,573 labeled entities, and 1,828,597 words. The entities have been annotated and classified into six categories, which are vulnerability type, affected products, root

Table 3: Overall Results of Our NER Model and Baselines

Performance	Baselines		Our Model
	en_core_web_sm	en_core_web_lg	
Precision	95.88	95.2	**98.75**
Recall	96.07	96.65	**98.55**
F1 Score	95.97	95.92	**98.65**

Table 4: Performance of Our NER Model on Individual Vulnerability Entities

Entity	Precision	Recall	F1 Score
Attacker Type	97.55	98.08	97.81
Impact	99.59	99.59	99.59
Attack Vector	98.48	97.6	98.04
Root Cause	98.68	99.02	98.85
Vulnerability Type	99.03	98.35	98.69
Affected product	98.21	98.36	98.28

cause, impact, attacker type, and attack vector. It is worth noting that these entities are in line with the vulnerability definitions of the official CVE template [31]. We split the dataset into training, validation, and test sets based on the 8:1:1 split ratio.

Baselines and Evaluation Metrics. Although our NER model encoder is constructed on a transformer encoder, we showcase its superior adaptability by constructing baseline NER models utilizing alternative encoders. Specifically, as our NER model is developed atop spaCy [21], we develop baselines utilizing spaCy's en_core_web_sm and en_core_web_lg encoders [21]. We train, fine-tune, and evaluate our NER model and baselines on the same training[1], validation, and test datasets. We evaluate our NER model and baselines with macro metrics, i.e., precision, recall, and F1 score [39].

Evaluation Performance. Table 3 presents the overall performance of the compared models across all the vulnerability entities. The results show the superior performance of our NER model compared to baselines. Moreover, we also evaluate our NER model on individual vulnerability entities. Table 4 shows the detailed performance, i.e., precision, recall, and F1 score, of our NER model on the individual vulnerability entities. Among the entities, our NER model performs the best on the "impact" entity but the worst on the "attacker type" entity. To understand the performance differences between these two entities, we conducted a manual investigation on our test samples and discovered that vulnerability descriptions often provide more explicit information about the impact rather than the attacker types. As an illustration, the NVD description of CVE-2017-11341 does not mention the attacker type in its description: "*There is a heap-based buffer over-read in lexer.hpp of LibSass 3.4.5. A crafted input will lead to a remote denial of service attack.*".

---

[1]The number of training steps is set to 2000, with which our NER model and baselines exhibited loss convergence.
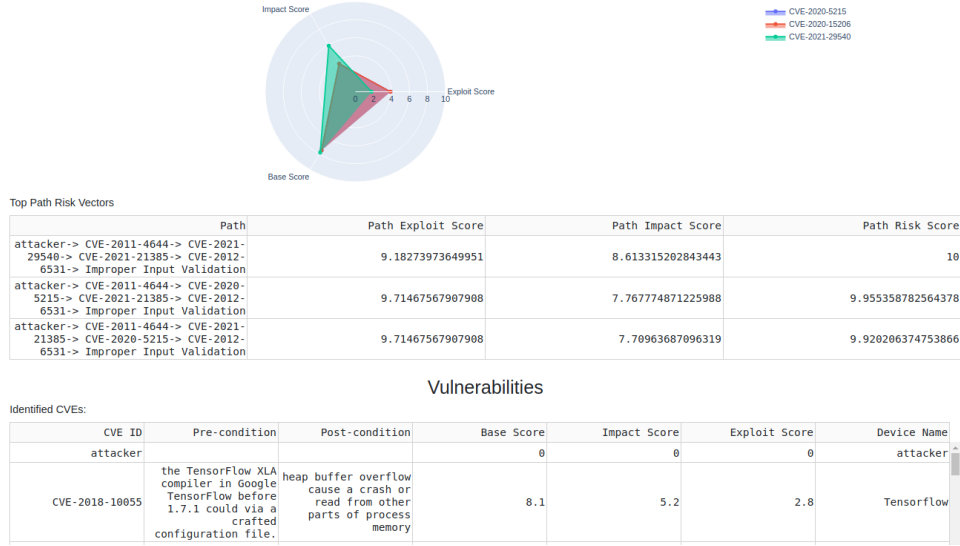
Figure 8: A small snapshot of GRAPHENE' front-end dashboard interface.
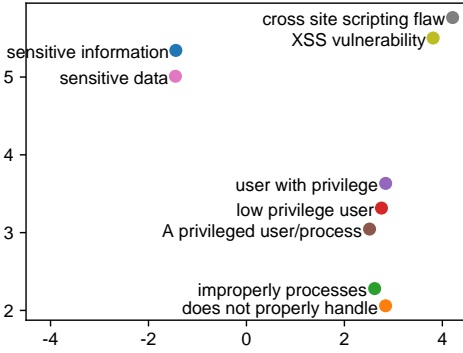

Figure 9: t-SNE Evaluation of Generated Embeddings. The distance between phrases shows their semantic similarity.

8.2 RQ2: Effectiveness of Our Word Embedding Method

To address **RQ2**, we perform evaluations on our word embedding methods for semantically matching attack graph nodes.

Word Embedding Dataset and Training. Based on the dataset curation framework presented in Figure 4, we are able to obtain 62,544 pre-processed security documents to train our word embedding models, which are stored in the SQLite database. The dataset sampling approach (§5.2) results in 6,335,336 word embedding training samples in total. To model the word semantics in the security context, we choose to train the Continuous Bag-of-Words (CBOW) and Skip-Gram models based on the SentencePiece framework [33]. We train the models to optimal performance by monitoring the loss curve which showed a convergence trend.

Evaluation. To evaluate word embedding, the common practice is using the public synonyms datasets, e.g., WS-353 and MTurk-287 [60]. However, to the best of our knowledge, there is no such open dataset in the security context. Additionally, constructing such a dataset with ground truth requires both too much human effort and domain-specific expertise, e.g., linguistic and security knowledge. Therefore, we opt to visualize word embeddings by using the t-distributed stochastic neighbor embedding (t-SNE) [59] to evaluate the performance of our word embedding models; t-SNE is an algorithm used for data visualization by reducing high-dimensional data to two or three dimensions [59]. It does so by preserving the local structure of the data and creating a low-dimensional map where similar data points are grouped together. We apply t-SNE to our word embedding models by visualizing the semantic similarity between attack circuit ports, e.g., the input of one attack graph node and the output of another attack graph node. Figure 9 presents such an example, where the distance of dots representing phrases corresponds to their semantic similarity. This figure shows that similar semantic phrases are clustered together, such as "cross site scripting flaw" and "XSS vulnerability". The results show that our word embedding models can capture the word semantics well in the security context.

8.3 Case Study

To evaluate GRAPHENE on capturing the holistic security posture of computer systems, we test it on the infrastructure shown in Figure 10. The test system includes three major components: (1) Jetson Nano device as the server with GPUs, (2) TP-link devices as network switch and router, and (3) Raspberry Pi boards as end devices. Moreover, we also adopt the common communication library, OpenSSL, as the software for securing network traffic.

To assess the security of the test system, we follow the workflow shown in Figure 6. That is, we first curate all 18K CVEs (till June 2022) from the NVD CVE database. We then identify CVEs related
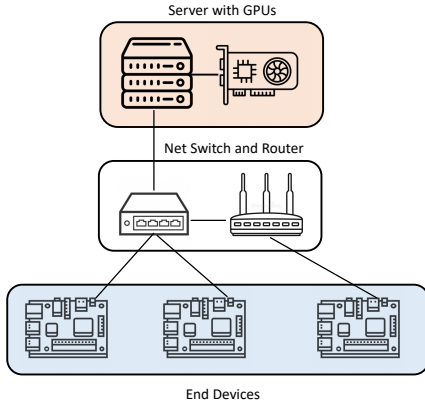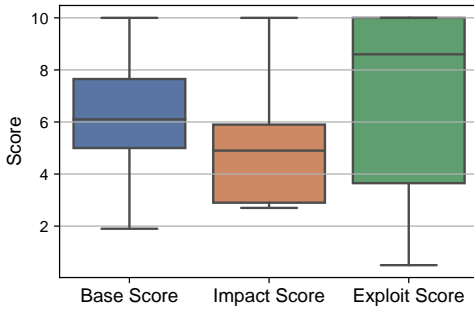
Figure 10: Test System for GRAPHENE

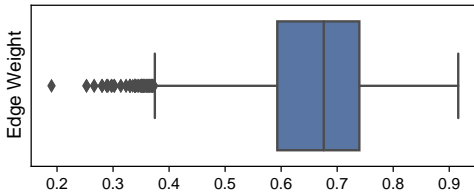Figure 11: Distribution of Vulnerability Scores

Figure 12: Distribution of Attack Graph Edge Weights

to the devices and library used in Figure 10 by matching the product names in CPE entries. After manually confirming the CVEs, we identified 99 CVEs vulnerability descriptions and reports for the test system, including 4 CVEs for Raspberry Pi, 8 CVEs for Jetson Nano, 35 CVEs for TP-link, and 52 CVEs for OpenSSL.

Next, we collect CVE metadata from the NVD database, i.e., CVE descriptions and vulnerability scores. CVE descriptions carry enriched security information. The mean, median, minimal, and maximum number of words in the descriptions are 48.5, 45, 9, and 131, respectively. Moreover, the most frequent bigram and trigram are "remote attackers" and "denial of service", which appear in the descriptions 38 and 34 times, showing the popular types of attackers and attacks. For risk assessment, Figure 11 presents the distribution of CVE base scores, impact scores, and exploit scores. From this figure, we observe that (1) CVSS base scores show a pretty high risk associated with the CVEs; (2) compared to base scores, the impact scores are in a narrow range; (3) the high impact scores reflect the high likelihood or probability that CVEs will be actively exploited in real-world attacks.

Attack Graph Construction. To identify attack graph nodes and build attack graph edges, we run REST APIs of the named entity recognition (NER) and word embedding models. The average inference time of the NER and word embedding models are 0.023 and 0.0017 seconds for each input, respectively. Afterward, we convert the identified entities into the input, output, precondition, and postcondition for each attack graph node. To build the edges, we calculate the Cosine similar score between attack graph nodes based on the embeddings generated by the word embedding models, which serve as the edge weights. Figure 12 presents the distribution of attack graph edge weights. The distribution shows that the majority of attack graph nodes are with high connectivity, i.e., the average edge weight is 0.683. In our experiments, we applied a threshold of 0.8 for the similarity score. In practical terms, if CVE *A* has a similarity score lower than 0.8 with CVE *B*, there is no edge connecting them in the resulting attack graphs. Our resulting cumulative attack graph contains 100 graph nodes, including one attacker node, 80 CVE nodes, and 19 target (CWE) nodes.

Risk Scoring.. Table 5 shows some numerical results as well as performance measures for the scenario illustrated in Figure 10. We provide the results for the cumulative attack graph as well as the layered attack graphs. As mentioned in Section 6, the exploit, impact, and risk scores are within the [0-10] range. The score computation time (expressed in seconds) includes the time needed to traverse the generated attack graph from the graph service and compute the graph exploitability, risk, and impact scores. The risk analysis time (expressed in seconds) is the time needed to rest the risk scoring system operations, that is, (1) identifying the shortest attack paths and the high-severity paths, (2) identifying key vulnerabilities, and (3) the minimum set of vulnerabilities that cover all the attack paths.

In the cumulative attack graph, GRAPHENE has identified CVE-2020-5215, CVE-2020-15206, and CVE-2021-29540 as the top three critical vulnerabilities associated with the installation of two vulnerable Tensorflow versions (2.4.2 and 1.15.2, respectively) on servers equipped with GPUs, as depicted in the spider graph in Figure 8. These vulnerabilities are notably part of the computed vertex cover, indicating their presence in the set of vulnerabilities encompassing all attack paths. Additionally, GRAPHENE highlighted the top three attack paths based on their elevated risk, exploitability, and impact, as shown in the "top path risk vectors" section of Figure 8. It is worth noting that in every generated attack graph, at least one node corresponds to the computed vertex cover, identifying the minimum set of vulnerabilities that must be addressed. Addressing these vulnerabilities through patching may render all attack graphs impractical. Ultimately, GRAPHENE has identified the five shortest attack paths within the cumulative attack graph, each having one CVE node with an exploitability score of 10, denoting the highest possible score. These CVE nodes are specifically CVE-2010-3173 (Mozilla Firefox on end devices), CVE-2011-0392 (Cisco TelePresence Recording Server on the server), CVE-2012-6531 (Zend Framework on end devices), CVE-2012-0884 (OpenSSL in all devices and the server), and CVE-2015-0763 (Cisco Unified MeetingPlace on end devices). Notably, analogous observations hold true for the network, system & hardware, machine learning, and cryptography layers.

Table 5: Risk scoring results for the scenario described in Figure 10.

Layer	Commulative	Network	System & HW	ML	Crypto
Exploit score (/10)	2.89	3.57	3.068	5.74	2.94
Impact score (/10)	3.07	3.18	3.8322	5.82	4.277
Risk score (/10)	2.31	2.72	3.9583	5.41	2.61
Total nodes	100	27	30	30	22
Number of attack paths	27297	137	53	6289	20
Shortest attack paths	5	3	4	3	2
Vertex cover size	50	19	15	21	3
Score computation time (seconds)	0.0198	0.0018	0.0014	0.0051	0.00098
Risk analysis time (seconds)	1.1874	0.0053	0.0021	0.2538	0.00097

9 RELATED WORK

Attack graph construction. For attack graph construction, prior approaches focused primarily on addressing the key challenge of extracting attack information [4, 25, 32, 36, 48, 54, 57, 62, 63]. These approaches can be categorized as rule-based and learning-based methods. For rule-based methods [4, 25, 48, 63], all rules are manually defined, limiting their scalability. Learning-based methods are built upon traditional machine learning algorithms [32, 62], such as support vector machine, or deep learning models [36, 54]. However, such learning-based methods are usually trained on fuzzy or domain-specific features but fail to learn the general vulnerability semantics in the security context. For example, these previous approaches have employed pre-trained word vectors that were trained on general English datasets. However, these methods tend to be less accurate in the security domain due to the presence of specific terminology and linguistic semantics that differ from those found in general English repositories.

Named Entity Recognition and Word Embedding for Security Applications. Natural language processing (NLP) has been widely used for security applications, e.g., binary analysis [12, 27] and security patch detection [64]. The security applications of NLP typically involve the utilization of named entity recognition (NER) and word embedding techniques. The application of NER to the descriptions and analysis of vulnerabilities is not new, and previous works have analyzed vulnerabilities from different aspects with domain-specific NER [8, 28, 62]. For example, Weerawardhana et al. [62] propose identifying security concepts in vulnerability information focusing on home computer security. Binyamini et al. [8] design a framework to model the attack techniques from security vulnerability descriptions, where the attack surfaces are the main focuses (e.g., software, hardware, operating systems, and network ports). However, there are several key limitations in those approaches, including (1) their scope is limited, e.g., home computer [62] and attack surfaces [8], and (2) they use limited machine learning models, such as LSTM networks [8] and convolutional neural networks [19].

Word embedding generates distributed representations of natural language words for efficient computations [43, 44]. Similar to its applications in natural language processing, existing research has adopted it to security applications [52]. For example, Shen et al. [52] model the cyberattack steps by temporal word embedding. Srivastava et al. [55] proposes enhancing security NLP models by word embeddings. Corizzo et al. [11] learn the raw data representations of network traffic for intrusion detection with word

embeddings. Unlike existing works, we focus on contextualizing the vulnerability-specific word embeddings by training models on massive vulnerability reports, i.e., CVE descriptions.

Risk Scoring. Li et al. [37] have proposed a cost/benefit analysis of the attack graphs. Administrators have to manually provide cost and benefit values for every node in the attack graph. A cost may relate to the adversary in terms of the resources that they need to possess, whereas a benefit is for the organization. Lu et al. [42] use the graph neural networks (GNNs) model to rank attack graphs. The authors use Google's PageRank web search algorithm [7] to infer each step's importance in the attack graph. GRAPHENE, however, extends the CVSS scoring system, which therefore aligns with the assessment and expertise of the CVE assigner. Idika et al. [24] have proposed attack graph-based security metrics to analyze attack graphs. However, their analysis is focused on the shortest paths and the number of paths an attacker can take to violate a security policy and not on the severity of the attack paths.

Other relevant approaches to security metrics are the probabilistic ones [3, 41, 61]. Abraham et al. [3] use stochastic modeling techniques based on Markov chains to analyze the network's current and future security state. Liu et al. [41] propose an approach based on Bayesian networks to model potential attack paths. However, for such an approach to work users have to provide the edge exploitability probability for each CVE. However, this approach may provide inaccurate analyses due to an over- or underestimated probability assignment. In addition, both of those approaches do not consider the attack's impact on the network infrastructure in terms of resource criticality.

10 CONCLUSION AND FUTURE WORK

The unceasing growth of cybercrime and network complexity require innovative cybersecurity solutions. GRAPHENE, the security analyzer proposed in this paper, is a promising step toward meeting this demand. Through the use of machine learning, natural language processing, and systematic vulnerability analysis, GRAPHENE offers a comprehensive security solution for general computing infrastructures in that it is capable of extracting semantic meanings from disclosed vulnerabilities, performing layered classification, and creating exhaustive directed attack graphs for holistic security analytics.

The distinctive multi-layered approach to security posture analysis provided by GRAPHENE holds the potential for a more nuanced understanding of a network's security landscape, allowing vulnerabilities of different natures and their interrelations to be effectively addressed. Furthermore, the proposed scoring methods provide a more tailored and relevant security assessment for the infrastructure under analysis, thereby enhancing the practical value of the analytics.

While the application of GRAPHENE shows promise, there is significant scope for future work. First, improving the machine learning models used in GRAPHENE will be crucial. As these models evolve, the accuracy and effectiveness of semantic extraction of vulnerabilities can be enhanced, leading to a more reliable and detailed vulnerability analysis. Second, enhancing the granularity of the layer classification process is an interesting research direction. While GRAPHENE classifies vulnerabilities into layers, refining this

classification to include sub-layers or specialized categories might provide more targeted insights for security measures. Third, developing advanced scoring methods can be a pivotal area of future exploration. As Graphene traverses the generated graph, devising innovative and more granular metrics to gauge the severity of vulnerabilities could make risk assessment more precise.

Furthermore, Graphene could be extended to incorporate automated patching and response capabilities to enhance its practical utility. This would allow for immediate remediation of identified vulnerabilities, enhancing the network's overall security. Integrating Graphene with real-time monitoring systems could also provide dynamic, up-to-date security posture analytics, thereby enabling faster response times to potential threats. Last, real-world testing of Graphene across various industries, like healthcare, energy, or finance, would provide invaluable feedback for fine-tuning the system. These tests could also uncover specific industry-related vulnerabilities, which could be incorporated into the model for a more sector-specific approach.

Overall, Graphene offers significant strides in cybersecurity. Still, its evolution is a continuous journey that will shape the contours of a safer and more secure digital landscape.

REFERENCES

[1] Dash documentation. [Online]. Available: https://dash.plotly.com/
[2] Flask documentation. [Online]. Available: https://flask.palletsprojects.com/en/3.0.x/
[3] S. Abraham and S. Nair, "Cyber security analytics: a stochastic model for security quantification using absorbing markov chains," *Journal of Communications*, vol. 9, no. 12, pp. 899–907, 2014.
[4] M. U. Aksu, K. Bicakci, M. H. Dilek, A. M. Ozbayoglu, and E. ı. Tatli, "Automated generation of attack graphs using nvd," in *Proceedings of the Eighth ACM Conference on Data and Application Security and Privacy*, 2018, pp. 135–142.
[5] R. Bar-Yehuda and S. Even, "A local-ratio theorem for approximating the weighted vertex cover problem," *Annals of Discrete Mathematics*, vol. 25, no. 27-46, p. 50, 1985.
[6] B. Bezawada, I. Ray, and K. Tiwary, "Agbuilder: an ai tool for automated attack graph building, analysis, and refinement," in *Data and Applications Security and Privacy XXXIII: 33rd Annual IFIP WG 11.3 Conference, DBSec 2019, Charleston, SC, USA, July 15–17, 2019, Proceedings 33.* Springer, 2019, pp. 23–42.
[7] M. Bianchini, M. Gori, and F. Scarselli, "Inside pagerank," *ACM Transactions on Internet Technology (TOIT)*, vol. 5, no. 1, pp. 92–128, 2005.
[8] H. Binyamini, R. Bitton, M. Inokuchi, T. Yagyu, Y. Elovici, and A. Shabtai, "A framework for modeling cyber attack techniques from security vulnerability descriptions," in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 2574–2583.
[9] H. Booth, D. Rike, and G. A. Witte, "The national vulnerability database (nvd): Overview," 2013.
[10] R. A. Bridges, C. L. Jones, M. D. Iannacone, K. M. Testa, and J. R. Goodall, "Automatic labeling for entity extraction in cyber security," *arXiv preprint arXiv:1308.4941*, 2013.
[11] R. Corizzo, E. Zdravevski, M. Russell, A. Vagliano, and N. Japkowicz, "Feature extraction based on word embedding models for intrusion detection in network traffic," *Journal of Surveillance, Security and Safety*, 2020.
[12] S. H. Ding, B. C. Fung, and P. Charland, "Asm2vec: Boosting static representation robustness for binary clone search against code obfuscation and compiler optimization," in *2019 IEEE Symposium on Security and Privacy (SP).* IEEE, 2019, pp. 472–489.
[13] "2022 official cybercrime report," 2022, https://www.esentire.com/resources/library/2022-official-cybercrime-report.
[14] J. Fan, Y. Li, S. Wang, and T. N. Nguyen, "Ac/c++ code vulnerability dataset with code changes and cve summaries," in *Proceedings of the 17th International Conference on Mining Software Repositories*, 2020, pp. 508–512.
[15] Z. Fang, H. Fu, T. Gu, P. Hu, J. Song, T. Jaeger, and P. Mohapatra, "Iota: A framework for analyzing system-level security of iots," in *2022 IEEE/ACM Seventh International Conference on Internet-of-Things Design and Implementation (IoTDI).* IEEE, 2022, pp. 143–155.
[16] X. Feng, X. Liao, X. Wang, H. Wang, Q. Li, K. Yang, H. Zhu, and L. Sun, "Understanding and securing device vulnerabilities through automated bug report

[17] analysis," in *SEC'19: Proceedings of the 28th USENIX Conference on Security Symposium*, 2019.
[17] FIRST.Org. Common vulnerability scoring system v3.1: Specification document. [Online]. Available: https://www.first.org/cvss/v3.1/specification-document
[18] R. Grishman and B. M. Sundheim, "Message understanding conference-6: A brief history," in *COLING 1996 Volume 1: The 16th International Conference on Computational Linguistics*, 1996.
[19] H. Guo, S. Chen, Z. Xing, X. Li, Y. Bai, and J. Sun, "Detecting and augmenting missing key aspects in vulnerability descriptions," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 31, no. 3, pp. 1–27, 2022.
[20] A. Hagberg, P. Swart, and D. S Chult, "Exploring network structure, dynamics, and function using networkx," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 2008.
[21] M. Honnibal, I. Montani, S. Van Landeghem, A. Boyd *et al.*, "spacy: Industrial-strength natural language processing in python," 2020.
[22] Y. Hu, I. Ameer, X. Zuo, X. Peng, Y. Zhou, Z. Li, Y. Li, J. Li, X. Jiang, and H. Xu, "Zero-shot clinical entity recognition using chatgpt," *arXiv preprint arXiv:2303.16416*, 2023.
[23] A. Ibrahim, S. Bozhinoski, and A. Pretschner, "Attack graph generation for microservice architecture," in *Proceedings of the 34th ACM/SIGAPP symposium on applied computing*, 2019, pp. 1235–1242.
[24] N. Idika and B. Bhargava, "Extending attack graph-based security metrics and aggregating their application," *IEEE Transactions on dependable and secure computing*, vol. 9, no. 1, pp. 75–85, 2010.
[25] M. Inokuchi, Y. Ohta, S. Kinoshita, T. Yagyu, O. Stan, R. Bitton, Y. Elovici, and A. Shabtai, "Design procedure of knowledge base for practical attack graph generation," in *Proceedings of the 2019 ACM Asia Conference on Computer and Communications Security*, 2019, pp. 594–601.
[26] X. Jin, C. Katsis, F. Sang, J. Sun, A. Kundu, and R. Kompella, "Edge security: Challenges and issues," *arXiv preprint arXiv:2206.07164*, 2022.
[27] X. Jin, J. Larson, W. Yang, and Z. Lin, "Binary code summarization: Benchmarking chatgpt/gpt-4 and other large language models," *arXiv preprint arXiv:2312.09601*, 2023.
[28] X. Jin, S. Manandhar, K. Kafle, Z. Lin, and A. Nadkarni, "Understanding iot security from a market-scale perspective," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1615–1629.
[29] X. Jin, K. Pei, J. Y. Won, and Z. Lin, "Symlm: Predicting function names in stripped binaries via context-sensitive execution-aware code embeddings," in *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022, pp. 1631–1645.
[30] Y. Jin, J. Lim, I. Yun, and T. Kim, "Compromising the macOS kernel through Safari by chaining six vulnerabilities," in *Black Hat USA Briefings (Black Hat USA)*, Las Vegas, NV, Aug. 2020.
[31] (2022) Key details phrasing. [Online]. Available: https://cveproject.github.io/docs/content/key-details-phrasing.pdf
[32] C. L. Jones, R. A. Bridges, K. M. Huffer, and J. R. Goodall, "Towards a relation extraction framework for cyber-security concepts," in *Proceedings of the 10th Annual Cyber and Information Security Research Conference*, 2015, pp. 1–4.
[33] T. Kudo and J. Richardson, "Sentencepiece: A simple and language independent subword tokenizer and detokenizer for neural text processing," *arXiv preprint arXiv:1808.06226*, 2018.
[34] M. Kuhlmann, C. Gómez-Rodríguez, and G. Satta, "Dynamic programming algorithms for transition-based dependency parsers," in *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, 2011, pp. 673–682.
[35] J. Li, A. Sun, J. Han, and C. Li, "A survey on deep learning for named entity recognition," *IEEE Transactions on Knowledge and Data Engineering*, vol. 34, no. 1, pp. 50–70, 2020.
[36] T. Li, Y. Guo, and A. Ju, "A self-attention-based approach for named entity recognition in cybersecurity," in *2019 15th International Conference on Computational Intelligence and Security (CIS).* IEEE, 2019, pp. 147–150.
[37] W. Li and R. B. Vaughn, "Cluster security research involving the modeling of network exploitations using exploitation graphs," in *Sixth IEEE International Symposium on Cluster Computing and the Grid (CCGRID'06)*, vol. 2. IEEE, 2006, pp. 26–26.
[38] Y. Li and T. Yang, "Word embedding for understanding natural language: a survey," in *Guide to big data applications.* Springer, 2018, pp. 83–104.
[39] P. Liu, Y. Guo, F. Wang, and G. Li, "Chinese named entity recognition: The state of the art," *Neurocomputing*, vol. 473, pp. 37–53, 2022.
[40] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
[41] Y. Liu and H. Man, "Network vulnerability assessment using bayesian networks," in *Data mining, intrusion detection, information assurance, and data networks security 2005*, vol. 5812. SPIE, 2005, pp. 61–71.
[42] L. Lu, R. Safavi-Naini, M. Hagenbuchner, W. Susilo, J. Horton, S. L. Yong, and A. C. Tsoi, "Ranking attack graphs with graph neural networks," in *Information Security Practice and Experience: 5th International Conference, ISPEC 2009 Xi'an,*

China, April 13-15, 2009 Proceedings 5. Springer, 2009, pp. 345–359.

[43] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.

[44] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," *Advances in neural information processing systems*, vol. 26, 2013.

[45] J. J. Miller, "Graph database applications and concepts with neo4j," in *Proceedings of the southern association for information systems conference, Atlanta, GA, USA*, vol. 2324, no. 36, 2013.

[46] D. Nadeau and S. Sekine, "A survey of named entity recognition and classification," *Lingvisticae Investigationes*, vol. 30, no. 1, pp. 3–26, 2007.

[47] M. Naili, A. H. Chaibi, and H. H. B. Ghezala, "Comparative study of word embedding methods in topic segmentation," *Procedia computer science*, vol. 112, pp. 340–349, 2017.

[48] J. Payne, K. Budhraja, and A. Kundu, "How secure is your iot network?" in *2019 IEEE International Congress on Internet of Things (ICIOT).* IEEE, 2019, pp. 181–188.

[49] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, "Scikit-learn: Machine learning in python," *the Journal of machine Learning research*, vol. 12, pp. 2825–2830, 2011.

[50] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1532–1543.

[51] R. Rehurek and P. Sojka, "Gensim–python framework for vector space modelling," *NLP Centre, Faculty of Informatics, Masaryk University, Brno, Czech Republic*, vol. 3, no. 2, 2011.

[52] Y. Shen and G. Stringhini, "{ATTACK2VEC}: Leveraging temporal word embeddings to understand the evolution of cyberattacks," in *28th USENIX Security Symposium (USENIX Security 19)*, 2019, pp. 905–921.

[53] P. M. Shishtla, K. Gali, P. Pingali, and V. Varma, "Experiments in telugu ner: A conditional random field approach," in *Proceedings of the IJCNLP-08 Workshop on Named Entity Recognition for South and South East Asian Languages*, 2008.

[54] K. Simran, S. Sriram, R. Vinayakumar, and K. Soman, "Deep learning approach for intelligent named entity recognition of cyber security," in *Advances in Signal Processing and Intelligent Recognition Systems: 5th International Symposium, SIRS 2019, Trivandrum, India, December 18–21, 2019, Revised Selected Papers 5.* Springer, 2020, pp. 163–172.

[55] S. Srivastava, B. Paul, and D. Gupta, "Study of word embeddings for enhanced cyber security named entity recognition," *Procedia Computer Science*, vol. 218, pp. 449–460, 2023.

[56] B. E. Strom, A. Applebaum, D. P. Miller, K. C. Nickels, A. G. Pennington, and C. B. Thomas, "Mitre att&ck: Design and philosophy," in *Technical report.* The MITRE Corporation, 2018.

[57] L. P. Swiler, C. Phillips, and T. Gaylor, "A graph-based network-vulnerability analysis system," Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), Tech. Rep., 1998.

[58] M. Urbanska, I. Ray, A. E. Howe, and M. Roberts, "Structuring a vulnerability description for comprehensive single system security analysis," *Rocky Mountain Celebration of Women in Computing, Fort Collins, CO, USA*, 2012.

[59] L. Van der Maaten and G. Hinton, "Visualizing data using t-sne." *Journal of machine learning research*, vol. 9, no. 11, 2008.

[60] B. Wang, A. Wang, F. Chen, Y. Wang, and C.-C. J. Kuo, "Evaluating word embedding models: Methods and experimental results," *APSIPA transactions on signal and information processing*, vol. 8, p. e19, 2019.

[61] L. Wang, T. Islam, T. Long, A. Singhal, and S. Jajodia, "An attack graph-based probabilistic security metric," in *Data and Applications Security XXII: 22nd Annual IFIP WG 11.3 Working Conference on Data and Applications Security London, UK, July 13-16, 2008 Proceedings 22.* Springer, 2008, pp. 283–296.

[62] S. Weerawardhana, S. Mukherjee, I. Ray, and A. Howe, "Automated extraction of vulnerability information for home computer security," in *International Symposium on Foundations and Practice of Security.* Springer, 2014, pp. 356–366.

[63] ——, "Automated extraction of vulnerability information for home computer security," in *Foundations and Practice of Security: 7th International Symposium, FPS 2014, Montreal, QC, Canada, November 3-5, 2014. Revised Selected Papers 7.* Springer, 2015, pp. 356–366.

[64] B. Wu, S. Liu, R. Feng, X. Xie, J. Siow, and S.-W. Lin, "Enhancing security patch identification by capturing structures in commits," *IEEE Transactions on Dependable and Secure Computing*, 2022.

[65] K. Xu, Z. Yang, P. Kang, Q. Wang, and W. Liu, "Document-level attention-based bilstm-crf incorporating disease dictionary for disease named entity recognition," *Computers in biology and medicine*, vol. 108, pp. 122–132, 2019.

[66] G. Zhou and J. Su, "Named entity recognition using an hmm-based chunk tagger," in *Proceedings of the 40th annual meeting of the association for computational linguistics*, 2002, pp. 473–480.

A KEYWORDS AND CWE LISTS FOR LAYER MATCHING

In this section, we list a subset of keywords used by Graphene to classify the vulnerabilities to one of our proposed layers; the Network Security layer. Table 6 lists the network-related keywords, whereas Table 7 lists a subset of widely adopted communication protocols in network infrastructures.

14

Table 6: The "Keywords" column corresponds to high-frequency keywords found in network vulnerability listings.

Keywords	Keywords
access control	MITM
authentication	nat
authenticity	NAT
authorization	network
availability	network interface
botnet	network packets
CDN	packets
certificate	port
certificates	ports
client	privacy
cloud	protocol
communication protocol	remote attacker
confidentiality	remote attackers
Cross-site request forgery	repudiation
Cross-site scripting	request
CSRF	response
DDoS	router
denial of service	sase
DoS	SDN
downgrade	server
edge network	side-channel
edge nodes	spoof
endpoints	spoofing
firewall	SQL
flood	switch
flooding	tamper
html	tampering
ICN	trust
injection	verification
input sanitization	VPN
input validation	wireless
integrity	XSS
IoT	zero-trust
LAN	ZTA
man-in-the-middle	link
message	network
mirai	

Table 7: The protocols are used for communication between entities.

Protocol	Protocol
tls	mqtt
ssl	coap
tcp	amqp
ip	lora
http	zigbee
https	WEP
ftp	WPA
ftps	icmp
udp	tor
lte	i2p
wifi	TELNET
bluetooth	DHCP
ARP	DNS

Table 8: The table shows a list of the related CWEs pertaining to network security vulnerabilities. The "CWE ID" is a unique weakness identifier, and "CWE Name" provides more information on the identifier.

CWE ID	CWE Name
20	Improper Input Validation
79	Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting')
80	Improper Neutralization of Script-Related HTML Tags in a Web Page (Basic XSS)
83	Improper Neutralization of Script in Attributes in a Web Page
87	Improper Neutralization of Alternate XSS Syntax
89	Improper Neutralization of Special Elements used in an SQL Command ('SQL Injection')
90	Improper Neutralization of Special Elements used in an LDAP Query ('LDAP Injection')
91	XML Injection (aka Blind XPath Injection)
93	Improper Neutralization of CRLF Sequences ('CRLF Injection')
97	Improper Neutralization of Server-Side Includes (SSI) Within a Web Page
98	Improper Control of Filename for Include/Require Statement in PHP Program ('PHP Remote File Inclusion')
113	Improper Neutralization of CRLF Sequences in HTTP Headers ('HTTP Request/Response Splitting')
183	Permissive List of Allowed Inputs
184	Incomplete List of Disallowed Inputs
200	Exposure of Sensitive Information to an Unauthorized Actor
209	Generation of Error Message Containing Sensitive Information
213	Exposure of Sensitive Information Due to Incompatible Policies
269	Improper Privilege Management
282	Improper Ownership Management
284	Improper Access Control
285	Improper Authorization
286	Incorrect User Management
287	Improper Authentication
287	Improper Authentication
288	Authentication Bypass Using an Alternate Path or Channel
289	Authentication Bypass by Alternate Name
290	Authentication Bypass by Spoofing
294	Authentication Bypass by Capture-replay
295	Improper Certificate Validation
296	Improper Following of a Certificate's Chain of Trust
297	Improper Validation of Certificate with Host Mismatch
298	Improper Validation of Certificate Expiration
299	Improper Check for Certificate Revocation
301	Reflection Attack in an Authentication Protocol
302	Authentication Bypass by Assumed-Immutable Data
303	Incorrect Implementation of Authentication Algorithm
304	Missing Critical Step in Authentication
305	Authentication Bypass by Primary Weakness
306	Missing Authentication for Critical Function
307	Improper Restriction of Excessive Authentication Attempts
308	Use of Single-factor Authentication
322	Key Exchange without Entity Authentication
345	Insufficient Verification of Data Authenticity
346	Origin Validation Error
352	Cross-Site Request Forgery (CSRF)
359	Exposure of Private Personal Information to an Unauthorized Actor
385	Covert Timing Channel
417	Communication Channel Errors
419	Unprotected Primary Channel
420	Unprotected Alternate Channel

Table 9: The table shows a list of the related CWEs pertaining to network security vulnerabilities. The "CWE ID" is a unique weakness identifier, and "CWE Name" provides more information on the identifier.

CWE ID	CWE Name
425	Direct Request ('Forced Browsing')
441	Unintended Proxy or Intermediary ('Confused Deputy')
497	Exposure of Sensitive System Information to an Unauthorized Control Sphere
515	Covert Storage Channel
522	Insufficiently Protected Credentials
564	SQL Injection: Hibernate
566	Authorization Bypass Through User-Controlled SQL Primary Key
593	Authentication Bypass: OpenSSL CTX Object Modified after SSL Objects are Created
599	Missing Validation of OpenSSL Certificate
601	URL Redirection to Untrusted Site ('Open Redirect')
603	Use of Client-Side Authentication
611	Improper Restriction of XML External Entity Reference
613	Insufficient Session Expiration
614	Sensitive Cookie in HTTPS Session Without 'Secure' Attribute
638	Not Using Complete Mediation
639	Authorization Bypass Through User-Controlled Key
643	Improper Neutralization of Data within XPath Expressions ('XPath Injection')
644	Improper Neutralization of HTTP Headers for Scripting Syntax
645	Overly Restrictive Account Lockout Mechanism
652	Improper Neutralization of Data within XQuery Expressions ('XQuery Injection')
706	Use of Incorrectly-Resolved Name or Reference
776	Improper Restriction of Recursive Entity References in DTDs ('XML Entity Expansion')
836	Use of Password Hash Instead of Password for Authentication
862	Missing Authorization
863	Incorrect Authorization
918	Server-Side Request Forgery (SSRF)
923	Improper Restriction of Communication Channel to Intended Endpoints
924	Improper Enforcement of Message Integrity During Transmission in a Communication Channel
939	Improper Authorization in Handler for Custom URL Scheme
940	Improper Verification of Source of a Communication Channel
941	Incorrectly Specified Destination in a Communication Channel
942	Permissive Cross-domain Policy with Untrusted Domains
1004	Sensitive Cookie Without 'HttpOnly' Flag
1211	Authentication Errors
1214	Data Integrity Issues
1220	Insufficient Granularity of Access Control
1263	Improper Physical Access Control
1270	Generation of Incorrect Security Tokens
1275	Sensitive Cookie with Improper SameSite Attribute
1311	Improper Translation of Security Attributes by Fabric Bridge
1327	Binding to an Unrestricted IP Address
1331	Improper Isolation of Shared Resources in Network On Chip (NoC)
1385	Missing Origin Validation in WebSockets