A Product-Oriented Assessment of Vulnerability Severity Through NVD CVSS Scores

Lucas Miranda, Lucas Senos, Daniel Menasché
Federal Univ. of Rio de Janeiro, RJ, Brazil

Gaurav Srivastava, Anton Kocheturov, Enrico Lovat, Abhishek Ramchandran
Siemens Corporate Technology, Princeton, NJ

Tobias Limmer
Siemens AG, Munchen, Germany

***Abstract*—Effective vulnerability assessment is critical in cybersecurity, especially for prioritizing risk mitigation in complex systems. The Common Vulnerability Scoring System (CVSS) provides a standardized method to evaluate the severity of vulnerabilities, but public repositories like the NVD often lack direct mappings between CVSS scores and specific affected products. This gap complicates practical decision-making for stakeholders responsible for remediation efforts. This paper presents a detailed analysis of how product-specific CVSS scores are assigned, focusing on the interplay between CVE Numbering Authorities (CNAs) and Common Platform Enumeration (CPE) vendors. We identify key challenges in aligning CVSS scores with specific products, especially for applications linked to multiple vendors, which often require additional context for accurate risk prioritization. These insights support the creation of automated tools to link CVSS scores to specific products and improve the transparency and consistency of vulnerability assessments. By addressing these gaps, this work provides practitioners with a framework to enhance vulnerability prioritization, particularly for products with diverse configurations and use cases.**

***Index Terms*—Cybersecurity, CVSS, CNA, vulnerability**

I. Introduction

The assessment of vulnerabilities in software and hardware products is one of the most fundamental and critical tasks in cybersecurity. The Common Vulnerability Scoring System (CVSS) serves as a widely adopted framework for evaluating the severity of vulnerabilities, identified by their Common Vulnerabilities and Exposures identifiers (CVEs), offering standardized metrics that assist stakeholders in prioritizing risk management efforts, e.g., product patch management. However, the process of assigning CVSS scores often involves multiple agents, including vendors, researchers, and other authorities, collectively known as CVE Numbering Authorities (CNAs). These CNAs are responsible for coordinating and issuing scores based on specific products, referred to by their Common Platform Enumeration (CPE) identifiers.

Why CNAs, CPEs and context matter. The topic of vulnerability severity has been extensively studied, with a wealth of research focused on assessing and standardizing metrics such as the CVSS. However, in practice, the products impacted by vulnerabilities and the products ultimately receiving patches are the central concern for practitioners. There is a significant missing link between vulnerability features, such as CVSS, and products, as the same vulnerability can affect multiple products in different ways, depending on their configurations, usage contexts, and underlying architectures

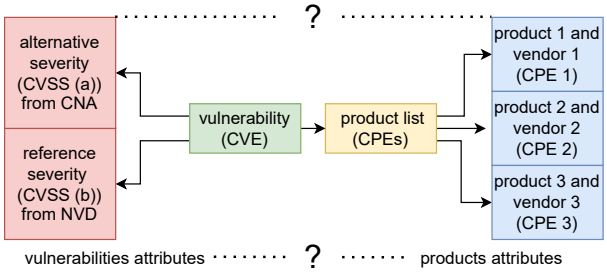

Fig. 1. How to prioritize patching of products? A single vulnerability may be related to multiple products and to multiple CVSS vectors. In general, it is challenging to determine product-specific CVSS vectors as there is no one-to-one matching between CVSS vectors and products – the number of products is typically much larger than the number of CVSS vectors for a given vulnerability.

(see Figure 1). This gap complicates the prioritization of remediation efforts and the application of effective patches. The goal of this paper is to begin bridging this gap by analyzing the relationship between vulnerabilities, affected products, and product-specific CVSS scores, providing insights that can enhance vulnerability management and improve the practical utility of CVSS assessments.

To improve vulnerability management, it is crucial to understand the relationship between CNAs and CPEs, as well as the distribution of CNA types across various products. Equally important is investigating the extent to which product-specific CVSS scores are issued, identifying the vendors most active in assigning these scores, and evaluating how much product-specific information is already embedded in vulnerability descriptions. With these goals in mind, this paper addresses the following research questions:

- **RQ1:** How are CNAs related to CPE vendors, and what patterns emerge in the assignment of CVSS scores?
  Motivation: Understanding the relationship between CNAs and CPE vendors (agents) helps identify who is assigning product-specific CVSS scores and why. For instance, when one single agent is responsible for all products associated with a given vulnerability, ambiguity about severity is reduced, improving vulnerability prioritization. Conversely, when there are multiple products from multiple agents related to the same CVE, it is challenging to determine how each product is impacted by that CVE.
- **RQ2:** Which vendors and products are most active in

979-8-3315-2096-0/25/$31.00 ©2025 IEEE

issuing product-specific CVSS scores?
Motivation: Identifying active vendors helps prioritize patching efforts and resources, focusing on vulnerabilities that affect widely used or critical systems.

- **RQ3:** What types of products (applications, hardware, operating systems) require product-specific CVSS scores, and how are they addressed?
Motivation: Exploring the types of products that require customized assessments highlights gaps in existing vulnerability management strategies and improves coverage for these categories.
- **RQ4:** How much relevant product-specific information is already conveyed in CVE descriptions, and how can it be better utilized?
Motivation: Leveraging existing CVE descriptions for semi-automatic CVSS assessment reduces manual overhead and improves consistency in vulnerability management practices.

The relationship between CNAs and products is complex because each CVE in the NVD is linked to at most one CNA different from NVD but can be associated with multiple CPEs, and not all CPEs are related to the issuing CNA (see Figure 2).

Our main contributions provide (partial) answers to the research questions under consideration.

- **RQ1**: We indicate that for 90% of vulnerabilities with product-specific CVSS scores, each CVE corresponds to a list of CPEs issued by a single CPE vendor.
Why it matters: When there is a single CPE vendor corresponding to a given vulnerability, the product patching prioritization is simplified, as in the simplest case all products from that vendor may be assumed to have the same priority and to correspond to the same CVSS vector.
- **RQ2**: We identify Microsoft, Cisco, and IBM as the top-3 most active CNAs.
Why it matters: When practitioners need to prioritize patching of products from these vendors, it is key to check product-specific CVSS vectors.
- **RQ3** and **RQ4**: We find that product-specific details for hardware and operating system products are easier to determine compared to applications, and that a significant portion of the information provided by CPEs is already present in CVE descriptions.
Why it matters: When performing patch prioritization, determining product-specific CVSS vectors for hardware and operating system modules is likely to require less effort than for general applications, which require further scrutiny to delve into details about the impact and exploitability of vulnerabilities in the context of interest.

By investigating these key areas, we aim to shed light on how CNAs and corresponding CPEs shape the landscape of vulnerability assessment, contributing to CNA and product-oriented vulnerability management strategies. This, in turn, complements previous efforts that focused on CVE-oriented [1]–[3] as opposed to product-oriented assessments.

This work is part of a research agenda on the analysis of product-specific scores [4], [5]. The key novel contributions of this paper on top of our previous works include an analysis of 1) the relation between CNAs and CPEs; 2) topmost CPE vendors, products and types; 3) the relation between vulnerability descriptions and CPEs.

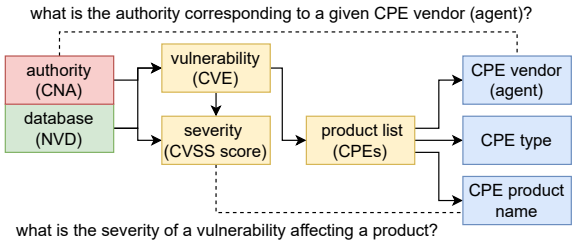

Fig. 2. NVD and CNAs provide public information about vulnerabilities (CVEs) and their severities (CVSS scores). A given CVE may correspond to multiple product-specific CVSS scores, one per CPE. However, there is a missing link between CVSS scores and corresponding products (CPEs).

Outline. The rest of this paper is organized as follows. Section II presents terminology. Our key findings are reported in Section III. Section IV presents related work and Section V concludes.

II. Terminology

Next, we define key terms that are crucial for understanding the concepts and analyses presented in this paper.

CVE (Common Vulnerabilities and Exposures): A unique identifier assigned to a specific security vulnerability.

CVSS (Common Vulnerability Scoring System): A standardized framework used to quantify the severity of a security vulnerability. CVSS scores range from 0 to 10.

Base Score: It is calculated based on factors including attack vector, attack complexity, privileges required, user interaction, and the impacts on confidentiality, integrity, and availability.

NVD (National Vulnerability Database): A U.S. government repository of publicly disclosed CVEs.

CNA (CVE Numbering Authority): CNAs are organizations authorized by MITRE to assign CVE IDs and provide CVSS scores. Organizations from around the world partner with the CVE Program as CNAs to assign CVE IDs and publish CVE Records for vulnerabilities within their distinct, agreed-upon scope. CNAs are classified into several types based on their role and nature: vendor, researcher, open source, CERT, hosted service, and bug bounty provider.[1] CNAs are responsible for numbering vulnerabilities that affect their products and reporting them to the NVD.

CPE (Common Platform Enumeration): A standardized method for naming software, hardware, and operating systems to facilitate automated vulnerability management. Each CPE corresponds to a product or set of products and is used to link vulnerabilities to specific products in the NVD.

CPE Agent: In this paper, we refer to the vendor appearing in the CPE string as the CPE agent. We do so to avoid

[1]CNA types from https://www.cve.org/PartnerInformation/ListofPartners.

TABLE I
CPE AGENTS AND CVES DISTRIBUTION (LEFT); CNA TYPES (RIGHT)

CPE agents	CVEs
1	31295
2	1856
3	902
4	294
5	140
≥6	70

CNA Type	# CNAs	# CVEs
Vendor	93	7658
Vendor, Open S.	29	2756
Vendor, Researcher	16	1522
Researcher	13	3943
CERT	9	889
Vendor, Researcher, Open Software	5	4930
Hosted Service, Open S., Vendor	Oracle	320
Open Source, Researcher	Snyk	538
Vendor et al.	Patchstack	2819
Other	12	169
Total	**180**	**25544**

confusion with the CNA vendor type discussed above. Note that the CPE agent can be a software vendor, an open-source project maintainer, or other stakeholders involved in product development and maintenance.

Product-Specific CVSS Scores: These are CVSS scores issued by CNAs for specific products, often influenced by unique configurations or environments. They may differ from the NVD's generalized CVSS scores, which apply across all products affected by the vulnerability.

Fuzzymatching: A technique used to identify similarities between entities (e.g., CNAs and CPE agents), based on string comparison algorithms. In the context of this paper, fuzzymatching is employed to associate product-specific vulnerabilities with CNAs based on similarity scores. Higher fuzzymatch scores indicate stronger confidence in the match.

III. EMPIRICAL FINDINGS

We collect data from NVD, on July 28, 2024. Our data comprises CVEs published from 2016 to 2024. We focus exclusively on vulnerabilities for which there are two CVSS scores issued by NVD and one additional distinct CNA. We analyze CVSS subscores, together with the CNAs reporting those scores. We found a total of 34,557 CVEs with two CVSS scores, regardless of whether they are equal or not. Out of those, 22,646 have at least one CVSS subscore that differs between the two sources. Except otherwise noted, we consider CVSS 3.x scores, as the fraction of vulnerabilities with CVSS 2.0 or 4.0 is negligible.

A. RQ1: How are CNAs related to CPE agents and what is the distribution of CNA types?

A clear majority of vulnerabilities exhibit a direct one-to-one correspondence between CNAs and CPE agents. Specifically, 90.6% of vulnerabilities involve a single CPE agent, as detailed in Table I (left). When a vulnerability is associated with a single CPE vendor, patching prioritization becomes simpler, as all products from that vendor can be assumed to share the same priority and CVSS vector.[2]

Table I (right) presents the distribution of CNAs by each type of CNA. Vendors represent the largest group, with 93 CNAs responsible for 7,658 CVEs, followed by combinations of vendors and open source or researcher roles. Independent researchers contributed 3,943 CVEs, while CERT organizations reported 889 CVEs. Notably, a small number of CNAs, such as those combining multiple roles (Vendor, Researcher, Open Source), are responsible for a significant portion of CVEs (4,930).

[2]In addition, when there is a single vendor, all relevant information about that vulnerability with respect to CVSS scores is likely to be found at that single vendor portal (e.g., Siemens portal for Siemens Security Advisories).

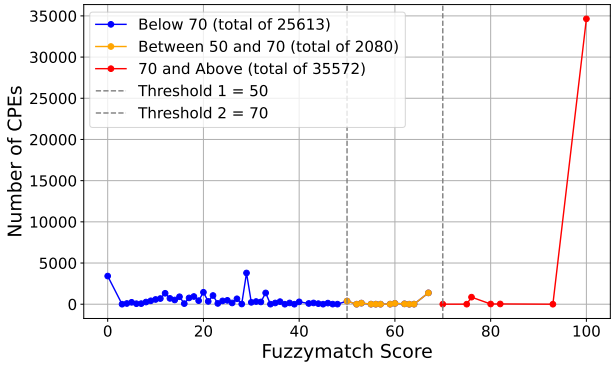

Fig. 3. Distribution of products (CPEs) with respect to fuzzymatch scores between their CPE agents and CNAs. Letting x be the score, when $x > 70$ there is a clear match between the CPE agent and the considered CNA, and for $50 \leq x < 70$ we manually inspected and vetted the proposed matchings.

Key takeaway: *For 90% of vulnerabilities with product-specific CVSS scores, there is a one-to-one mapping between CNAs and CPE agents, simplifying data retrieval for practitioners by centralizing relevant information in a single portal, such as that maintained by a CPE vendor.*

B. RQ2: What are the top-10 CPE agents issuing product-specific CVSS scores?

As pointed out in Figure 2, the relationship between CNAs and products is non-trivial. This is because each CVE at NVD has reference to at most one CNA but to multiple CPEs. Only some of those CPEs are related to the CNA that issued the CVSS. In this section, we aim to determine, for each (CVE, CNA) pair, which CPEs they correspond to. To this aim, we rely on fuzzymatching. After presenting the fuzzymatching results (Figure 3), we illustrate how they impact statistics about the top-10 CPE agents (Figure 4).

Figure 3 presents the distribution of CPEs based on their fuzzymatching scores when attempting to match a CPE agent (i.e., the vendor in the CPE string) against a CNA. Each CPE contains a CPE agent, product type, product name and version. The x-axis represents fuzzymatch scores ranging from 0 to 100, while the y-axis indicates the number of CPEs corresponding to each score range. The matchings are categorized into three distinct regions. For matchings with scores below 50, these were ignored due to low reliability. In the range between 50 and 70, which included 2,080 CPEs, the matching quality varied, and these matchings were manually inspected to vet inconclusive results. For matchings with scores of 70 and above, all matchings were assumed to be valid, after sampling a fraction of those matchings and verified their validity.

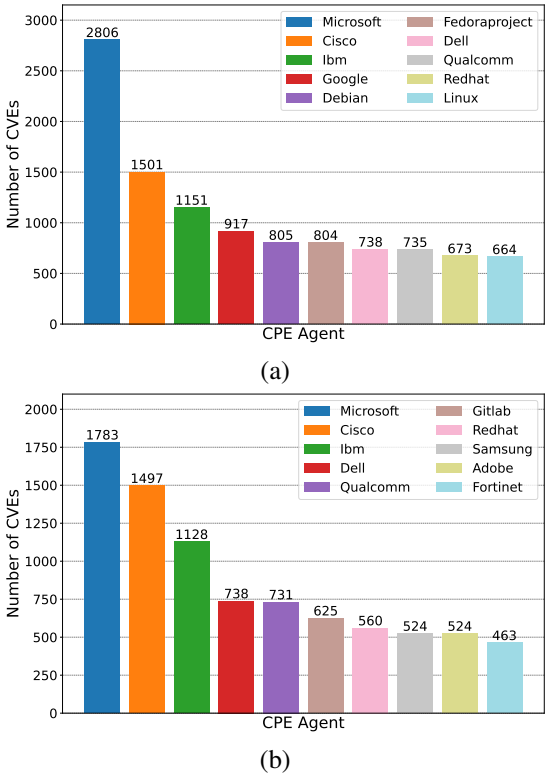

Fig. 4. Top-10 CPE agents: number of CVEs per CPE agent (a) without accounting for CNA matchings and (b) accounting for CNA matchings.

Figures 4(a) and 4(b) show the most relevant CPE agents by considering the total population of matchings (top) and those that matched against a CNA (bottom). In the first, vendors such as Microsoft, Cisco, IBM, Google, and Debian have the largest numbers of CVEs. Microsoft leads significantly, followed by Cisco and IBM. The second figure focuses on the subset of CPE agents whose fuzzymatching scores were high enough to match against a CNA, revealing a similar dominance by Microsoft, Cisco, and IBM. However, some differences are evident: vendors like Dell and Qualcomm play a more prominent role in this subset, while Google and Debian are less represented compared to the full population. This indicates that some CPE agents, like Microsoft and Cisco, maintain a strong presence in both the total population and matchings. Others, such as Dell and Qualcomm, appear in the top-10 list of CPE agents when accounting for the matchings between CPEs and CNAs, but do not appear in the top-10 list when matchings are not taken into account.

Key takeaway: *If a vulnerability has either its CNA or one of its CPE agents listed in Figure 4, practitioners should verify whether a product-specific CVSS vector is available, as this can provide critical context for prioritizing mitigation efforts.*

C. RQ3: Which products need product-specific CVSS scores?

To address this question, we examine three key types of products: applications, hardware, and operating systems. As shown in Figure 5, hardware and operating systems have the largest number of associated products, with 16,503 and 14,076

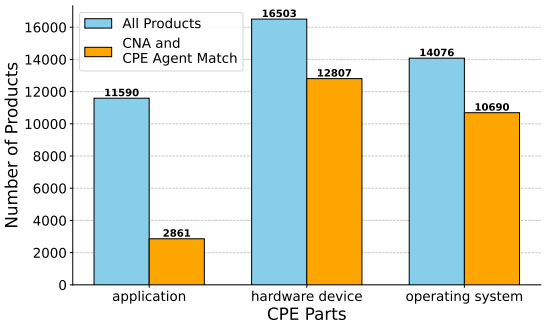

Fig. 5. Distribution of products across CPE parts categories. The blue bars represent all products, while the orange bars show only products where the CNA and CPE agents match.

cases, respectively. Applications account for 11,590 cases. When filtering for products whose CPE agent matched against CNAs, these numbers slightly drop, except for applications, which show a smaller count of 2,861 cases. These findings suggest two key points: 1) application products are often associated with multiple vendors, and NVD's limit of one additional CNA per CVE may be too restrictive; and 2) while hardware devices and operating systems tend to have product-specific CVSS scores at NVD, many applications may need alternative sources to find their corresponding CVSS scores.

Key takeaway: *Applications, hardware devices, and operating systems frequently require product-specific CVSS scores. However, applications are often linked to multiple CPE agents, unlike hardware and operating systems, which can complicate mitigation prioritization for practitioners.*

Proposed solution: hierarchical CVE identifier. *The current CVE ID structure is a flat numeric format. Introducing an optional hierarchical suffix to the CVE ID, referencing the associated CPE agent, could enable the assignment of distinct CVSS vectors to different CVE IDs. This would streamline product-specific mitigation prioritization for practitioners.*

D. RQ4: How much information about products and vendors is already conveyed in the vulnerability descriptions?

Answering this question is key for semi-automatic CVSS assessments using tools such as CVSS-BERT and to determine if NVD can be considered self-contained. Fig. 6 shows the Complementary Cumulative Distribution Function (CCDF) between agents, products and versions of CPE strings and CVE descriptions. As shown in Figure 6, the fuzzymatch between product names in vulnerability descriptions and those in CPE entries is notably easier than matching CPE agents or product versions. Product names achieve a match rate of 0.84, while CPE agents and versions exhibit lower match rates of 0.63 and 0.29, respectively, accounting for a fuzzymatch threshold of 70. This suggests that product-related information is well-conveyed in the vulnerability descriptions at NVD, though matching beyond product names can present challenges.

Key takeaway: *For practitioners, product information is often well-represented in NVD vulnerability descriptions, with product names being relatively easy to match. However, lower*

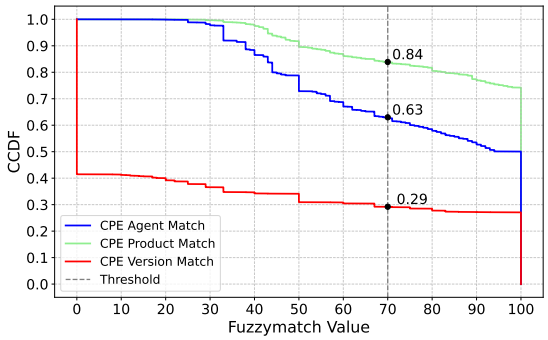

Fig. 6. CCDF of fuzzymatch scores between agents, products and versions of CPE strings and CVE description.

match rates for CPE agents and product versions highlight potential gaps, suggesting the need for supplementary tools or methods to enhance product-specific mitigation strategies.

IV. Related Work

Several studies have examined the consistency and accuracy of vulnerability assessments, particularly in the context of CVSS scores. [2], [3] compared the accuracy of assessments by students and security professionals, highlighting that both groups converge on similar outcomes, despite differences in expertise. In contrast, [1] revealed inconsistencies in CVSS evaluations, with our research providing insights framing differences in the realm of product-specific CVSS assessments.

Efforts to automate the matching of unstructured data, such as security advisories, to structured formats like CPEs have been explored in [6], [7], using string similarity techniques and customized metrics. This complements our goal of semi-automating product-specific CVSS score assessments with CPEs, an area not addressed by previous works.

Inconsistencies in databases like NVD have also been identified in prior research [8]–[12], though they focus on unintentional errors. Our work investigates intentional variations in CVSS scores due to CNA product-specific assessments.

While there is a growing body of literature on explaining and predicting CVSS scores based on vulnerability descriptions [13]–[19], none account for the differing assessments by CNAs. We propose that models from prior work [14], [17] can be fine-tuned for specific CNA needs, which we leave as subject for future exploration.

V. Conclusion

Vulnerability severity scores, such as CVSS scores, are designed to inform stakeholders, including end-customers and intermediate suppliers, about vulnerabilities affecting the products they use. It is essential for customers to identify the appropriate CVSS score for prioritizing mitigation efforts based on the specific product in use.

However, public databases like the NVD currently lack explicit connections between CVSS scores and individual products. While applications, hardware devices, and operating systems often require product-specific CVSS scores, our analysis shows that determining these scores is particularly complex for applications due to their frequent association with multiple vendors. Addressing this complexity is critical for enabling accurate vulnerability management and enhancing the effectiveness of semi-automated assessments.

Future research involves incorporating product-specific data into CVSS scoring mechanisms and establishing clearer links between vulnerability severity and affected products within public databases like the NVD. Such advancements would facilitate more transparent and explainable CVSS assessments, ultimately improving the prioritization and mitigation processes for practitioners.

References

[1] J. Wunder, A. Kurtz, C. Eichenmüller, F. Gassmann, and Z. Benenson, "Shedding Light on CVSS Scoring Inconsistencies: A User-Centric Study on Evaluating Widespread Security Vulnerabilities," in *IEEE Symposium on Security and Privacy*, 2023, pp. 58–58.

[2] F. Massacci, "The holy grail of vulnerability predictions," *IEEE Security & Privacy*, vol. 22, no. 1, pp. 4–6, 2024.

[3] L. Allodi, M. Cremonini, F. Massacci *et al.*, "Measuring the accuracy of software vulnerability assessments: experiments with students and professionals," *Empirical Softw. Engin.*, vol. 25, pp. 1063–1094, 2020.

[4] L. Miranda, L. Senos, D. Menasché, G. Srivastava, A. Kocheturov, A. Ramchandran, E. Lovat, and T. Limmer, "Learning CNA-Oriented CVSS Scores," in *CloudNet*. Rio de Janeiro, Brazil: IEEE, 2024.

[5] L. S. Coutinho, L. G. Miranda, D. S. Menasché, G. Srivastava, A. Kocheturov, A. Ramchandran, E. Lovat, and T. Limmer, "How Context Impacts Vulnerability Severity? An Analysis of Product-Specific CVSS Scores," in *LADC*, Recife, Brazil, 2024.

[6] K. McClanahan and Q. Li, "Towards Automatically Matching Security Advisories to CPEs: String Similarity-based Vendor Matching," in *ICNC Workshop*, 2024.

[7] K. McClanahan, S. Elder, M. L. Uwibambe, Y. Liu, R. Heng, and Q. Li, "When chatgpt meets vulnerability management: the good, the bad, and the ugly," in *IEEE Int'l Conf. on Computing, Networking and Communications (ICNC)*, 2024.

[8] R. Croft *et al.*, "An investigation into inconsistency of software vulnerability severity across data sources," in *SANER*, 2022, pp. 338–348.

[9] S. Zhang, M. Cai, M. Zhang, L. Zhao, and X. d. C. de Carnavalet, "The Flaw Within: Identifying CVSS Score Discrepancies in the NVD," in *CloudCom*. IEEE, 2023, pp. 185–192.

[10] A. Anwar *et al.*, "Cleaning the NVD: Comprehensive quality assessment, improvements, and analyses," *IEEE Transactions on Dependable and Secure Computing*, vol. 19, no. 6, pp. 4255–4269, 2021.

[11] Y. Dong, W. Guo, Y. Chen, X. Xing, Y. Zhang, and G. Wang, "Towards the detection of inconsistencies in public security vulnerability reports," in *USENIX Security*, 2019, pp. 869–885.

[12] R. Croft, M. A. Babar, and M. M. Kholoosi, "Data quality for software vulnerability datasets," in *2023 IEEE/ACM 45th International Conference on Software Engineering (ICSE)*. IEEE, 2023, pp. 121–133.

[13] R. Ghosh, O. Farri *et al.*, "CVE-LLM: Automatic vulnerability evaluation in medical device industry using large language models," *arXiv preprint arXiv:2407.14640*, 2024.

[14] C. Elbaz, L. Rilling, and C. Morin, "Fighting n-day vulnerabilities with automated cvss vector prediction at disclosure," in *Int. Conference on Availability, Reliability and Security*, 2020, pp. 1–10.

[15] J. C. Costa, T. Roxo, J. B. Sequeiros, H. Proenca, and P. R. Inacio, "Predicting CVSS metric via description interpretation," *IEEE Access*, vol. 10, pp. 59 125–59 134, 2022.

[16] P. Kühn, D. N. Relke, and C. Reuter, "Common vulnerability scoring system prediction based on open source intelligence information sources," *Computers & Security*, vol. 131, p. 103286, 2023.

[17] M. R. Shahid and H. Debar, "CVSS-BERT: Explainable natural language processing to determine the severity of a computer security vulnerability from its description," in *ICMLA*. IEEE, 2021, pp. 1600–1607.

[18] Z. Han *et al.*, "Learning to predict severity of software vulnerability using only vulnerability description," in *ICSME*. IEEE, 2017.

[19] A. Khazaei, M. Ghasemzadeh, and V. Derhami, "An automatic method for CVSS score prediction using vulnerabilities description," *Journal of Intelligent & Fuzzy Systems*, vol. 30, no. 1, pp. 89–96, 2016.