

CVE-LLM : Ontology-Assisted Automatic Vulnerability Evaluation Using Large Language Models

Rikhiya Ghosh¹, Hans-Martin von Stockhausen², Martin Schmitt²,
George Marica Vasile³, Sanjeev Kumar Karn¹, Oladimeji Farri¹

¹Digital Technology and Innovation, Siemens Healthineers USA

²Cybersecurity, Siemens Healthineers Germany

³Corporate Technology, Siemens AG Romania
rikhiya.ghosh@siemens-healthineers.com

Abstract

The National Vulnerability Database (NVD) publishes over a thousand new vulnerabilities monthly, with a projected 25 percent increase in 2024, highlighting the crucial need for rapid vulnerability identification to mitigate cybersecurity attacks and save costs and resources. In this work, we propose using large language models (LLMs) to learn vulnerability evaluation from historical assessments of medical device vulnerabilities in a single manufacturer’s portfolio. We highlight the effectiveness and challenges of using LLMs for automatic vulnerability evaluation and introduce a method to enrich historical data with cybersecurity ontologies, enabling the system to understand new vulnerabilities without retraining the LLM. Our LLM system integrates with the in-house application - Cybersecurity Management System (CSMS) - to help Siemens Healthineers (SHS) product cybersecurity experts efficiently assess the vulnerabilities in our products. Also, we present guidelines for efficient integration of LLMs into the cybersecurity tool.

Introduction

A software vulnerability is a flaw that malicious attackers can exploit to compromise a system’s confidentiality, integrity or availability (CIA), such as buffer overflows, SQL injection, and XSS. Regulatory authorities, non-profits, and companies have created open-source standards for sharing vulnerability information, resulting in continually updated databases by cybersecurity professionals. Examples of such databases include: (a) Common Vulnerabilities and Exposures (CVE)¹: Catalogs software vulnerabilities (Mitre 2005), (b) Common Weakness Enumeration (CWE)²: Describes software security weaknesses, (c) Common Attack Pattern Enumeration and Classification (CAPEC)³: Details known attack and mitigation patterns, and (d) Common Vulnerability Scoring System (CVSS): Assesses vulnerability severity (Mell, Scarfone, and Romanosky 2006).

Among these, the CVE database stands out for its detailed descriptions and assessments of vulnerabilities. The

CVE description offers a concise overview of the vulnerability, outlining its nature, affected products, potential impact, and conditions for exploitation. CVSS reports CVE severity in two forms: CVSS Score and CVSS Vector. The CVSS Score, a composite score on a scale of 1 to 10, offers a quick overview of the vulnerability’s impact. It is derived from CVSS Vectors, which comprise several metrics evaluating the exploitability of the vulnerability, such as Base (inherent characteristics), Temporal (changes over time), and Environmental (local CIA requirements and controls). The CVE database includes Base and Temporal metrics, while Environmental metrics are calculated by cybersecurity experts for specific products, ensuring accurate severity assessment within their context.

Medical devices are particularly vulnerable to security threats that can compromise patient safety, privacy, medical data integrity, and the availability of diagnostic or treatment devices. According to the US FDA guidance⁴, Medical Device Manufacturers (MDMs) must monitor third-party software components for emerging vulnerabilities and assess their impact on the device’s safety and security.

Vulnerability management involves three steps: detection, evaluation, and mitigation. Component vendors detect and issue notifications for vulnerabilities, which MDMs must monitor and assess for relevance to their products. Mitigation may involve vendor patches or MDM-implemented controls. Sometimes, CVE vulnerabilities may not pose a risk, requiring only communication to customers.

The notifications generated by component vendors consist of one or more CVEs that impact their component. When these notifications are issued, Siemens Healthineers (SHS) Cybersecurity experts conduct manual vulnerability evaluations utilizing CVE descriptions, CVSS scores and the related knowledge in CWE and CAPEC. In SHS, vulnerability evaluation involves the following assessments based on asset details and the notification : (a) VEXCATEGORY, (b) VEXJUSTIFICATION, (c) INTERNALCOMMENT, (d) CUSTOMERCOMMENT, and (e) VECTOR. Table 1 provides a brief description of these evaluation types.

Given the extensive product portfolios ($\approx 1.7K$ prod-

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹<https://www.mitre.org/>

²<https://cwe.mitre.org/>

³<https://capec.mitre.org/>

⁴Food and Drug Administration (FDA) Guidance on Postmarket Management of Cybersecurity in Medical Devices : <https://www.fda.gov/media/95862/download>

| EVALUATIONS | Description | Value Type |
|------------------|--|---|
| VEXCATEGORY | Indicates whether the asset is affected by the notification. | Binary categorical values |
| VEXJUSTIFICATION | Provides further explanation if the asset is not affected. | Multiclass Categorical values |
| INTERNALCOMMENT | Details the problem and advised internal solution. | Text |
| CUSTOMERCOMMENT | Summarizes the notification’s impact and solution directed towards the customer. | Text |
| VECTOR | Evaluates the CVSS Environmental metrics. | Multiclass, Multilabel categorical values |

Table 1: Details for EVALUATIONS performed manually currently by SHS cybersecurity experts

ucts) of SHS and numerous third-party components, manually performing these evaluations (≈ 1.7 million/month) can be overwhelming. Automation of vulnerability evaluations serves two purposes: (a) quick and efficient management of the large volume of evaluations required, ensuring consistent quality over time, and (b) mitigation of the risk associated with the potential unavailability of knowledgeable experts, who may no longer be in the same role or company during the entire post-market lifetime of a device, which can span up to 20 years.

From a machine learning perspective, SHS vulnerability evaluation tasks in Table 1 could be categorized into two primary tasks: text generation (INTERNALCOMMENT and CUSTOMERCOMMENT) and text classification (VEXCATEGORY, VEXJUSTIFICATION, and VECTOR). Vector generation is a multi-label, multi-class classification problem, VEXCATEGORY is binary classification, and VEXJUSTIFICATION is multi-class classification.

Recent advancements in transformer models (Devlin et al. 2018) have significantly improved language generation and reasoning capabilities across various natural language processing tasks (Wu et al. 2023; Li et al. 2023; Yang et al. 2024; Sallam 2023). While encoder models have traditionally been employed for CVE classification in vulnerability evaluation, the progress in generative AI, particularly decoder-based LLMs, has largely been applied only to penetration testing (Happe and Cito 2023; Deng et al. 2024). There is a lack of real-world experiments on the effectiveness of LLMs in vulnerability evaluation. To address the challenge MDMs face in assessing a large number of vulnerabilities, we present an LLM trained on historical evaluation data to assess the impact of vulnerabilities on assets. This approach enables faster detection and assessment of exploitable vulnerabilities, leading to quicker mitigation.

In this work, we developed and deployed an LLM to assist product cybersecurity experts in vulnerability evaluation. Overall our contributions are:

1. To the best of our knowledge, this is the first work that explores using LLMs for vulnerability assessment based solely on asset and vulnerability descriptions in medical technology industry setting.
2. We benchmarked our model against other fine-tuned open source LLMs and provide insights into best practices for training and deploying a vulnerability LLM.
3. Additionally, we demonstrated how incorporating knowl-

edge from vulnerability knowledge graphs into LLM training enhances performance, particularly in deployment settings.

Datasets

Our methodology adheres to the pretrain-then-finetune paradigm. We conduct Domain Adaptive Pretraining (DAPT) using the open-source MPT-7B model (Team 2023), which is followed by Supervised Finetuning (SFT) to develop our model, designated as CVE-LLM. Thus, we construct two distinct datasets: the DAPT dataset for the pre-training phase and the SFT dataset for the subsequent fine-tuning phase.

DAPT Dataset

We have formed a DAPT dataset consisting of 350K vulnerability description documents, which combine publicly available CVEs (248K) from the NVD (National Vulnerability Database) and SHS organization-wide vulnerability documents (102K). The organization-wide vulnerability documents describe a combination of one or more vulnerabilities,

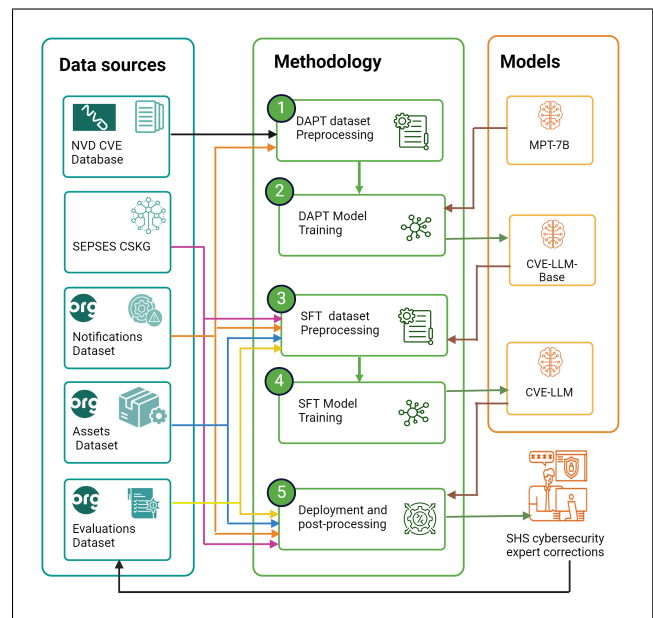


Figure 1: Schematic overview of CVE-LLM

| Evaluation | Instruction | Response |
|------------------------------|---------------------------------|--|
| VEXCATEGORY/VEXJUSTIFICATION | What is the category? | <i>Justification. Category: Category</i> |
| INTERNALCOMMENT | Generate internal comment. | <i>Internal Comment</i> |
| CUSTOMERCOMMENT | Generate customer comment. | <i>Customer Comment</i> |
| VECTOR | Generate environmental vectors. | <i>Environmental vector description</i> |

Table 4: Instructions for each evaluation type

While the CVE entries describe particular instances of vulnerabilities, their corresponding CWE entries describe general types of weaknesses that can lead to these vulnerabilities. The CAPEC database, on the other hand, provides actionable insights into how those weaknesses are exploited through specific attack techniques and patterns. In the CAPEC database, prerequisites refer to the specific conditions or requirements that must be met for an attack pattern to be successfully executed. Typical severity refers to the level of impact or potential damage that an attack pattern can typically inflict if successfully executed. Mitigations refer to strategies, techniques, or practices that can be employed to prevent or reduce the effectiveness of the attack patterns described.

We utilize the SEPSES knowledge graph to identify the CAPEC entries associated with the linked CWE entries relevant to the CVEs included in a notification to enrich the notification description.

Formation of SFT Dataset The format of the SFT dataset, as shown in Table 2, uses the standard Alpaca format (Taori et al. 2023). Each evaluation in the EVALUATIONS dataset forms four entries in the SFT dataset. Table 4 defines the instruction for each evaluation type. An entry in the SFT dataset contains the corresponding instruction, a description of the asset from ASSETS dataset, the ontology-enriched description of the notification from NOTIFICATIONS dataset, the affected software components, CVSS base and temporal metric descriptions, and one of the corresponding assessments: VEXCATEGORY (and VEXJUSTIFICATION), INTERNALCOMMENT, CUSTOMERCOMMENT, and CVSS VECTOR description. CAPEC mitigations are added only for the generation of Internal Comment, while prerequisites and typical severity are added for all evaluation types. The SFT dataset finally goes through text cleaning, removal of incomplete evaluations and long texts (number of tokens >1024).

Model Training

DAPT Model Training The DAPT dataset is randomly split into 90:10 split with 315K training data points and 35K validation data points. We have continuously pretrained MPT-7B base model autoregressively for next token prediction with an objective function of cross entropy. We have expanded the vocabulary of the MPT-7B base model to include names of the components and the organization software, and added 539 new tokens to the 50K vocabulary of MPT. The model is trained using DeepSpeed (Rasley et al. 2020) zero-3 optimization with Lion optimizer (Chen et al.

2024) (learning rate: 1e-4, weight decay: 1e-2) for 3 epochs on 8 Tesla A100 SXM4 GPUs.

SFT Model Training The SFT dataset has 750K unique instructions, and the dataset is randomly split into 600K training, 75K validation, and 44K test data points. The DAPT model is finetuned with the training and validation data by back propagating on completions only, while setting the token labels of the instructions to -100. Except for the learning rate (1e-5), similar optimization, hyperparameters and GPU resources for training are used for instruction-tuning.

Deployment

The deployment framework includes generation of evaluations using the CVE-LLM, using pre-defined rules to avoid easily identifiable erroneous generations, and using faster inference frameworks to improve model throughput. The AI-generated vulnerability evaluations are validated by a product cybersecurity expert (with higher priority given to the evaluations with Affected VEXCATEGORY). The corrected human evaluations form the input for subsequent retraining phase.

Generation of Evaluations Evaluation generation follows a zero shot model inference where the model is provided a prompt in the pre-defined SFT instruction format, containing information about sub-organization name, software name and version, asset name and version, notification description, CAPEC entries, list of component descriptions for components present in both asset and notification, base and temporal versions and CVSS version. The inference/test dataset is divided into two sub-datasets based on the size of tokens in the inference instruction. For instructions longer than 920 tokens, the sequence length of the trained model is increased to 150 + maximum token length of the longest instruction. If the size of the instruction is less than 920 tokens, the trained model without any changes is used to inference the dataset.

Rule-Based Correction Post-processing of model outputs involves using CSMS-specific knowledge in the form of pre-defined rules to correct mistakes in model generation. The rules are as follows:

1. If VEXCATEGORY value is NotAffected, the evaluation vector is not generated.
2. If VEXJUSTIFICATION generated is any text other than the pre-defined categories, 'Other' gets assigned to VEXJUSTIFICATION.

| Evaluation | Llama3-RAG | Mistral-7B | Llama2-7B | CVE-LLM | CVE-LLM-Eval | CVE-LLM-Prod |
|------------------|------------|------------|-----------|---------|--------------|--------------|
| VEXCATEGORY | 0.25 | 0.64 | 0.56 | 0.94 | 0.94 | 0.86 |
| VEXJUSTIFICATION | 0.12 | 0.60 | 0.53 | 0.90 | 0.92 | 0.95 |
| INTERNALCOMMENT | 0.19 | 0.71 | 0.62 | 0.79 | 0.79 | 0.73 |
| CUSTOMERCOMMENT | 0.28 | 0.70 | 0.64 | 0.88 | 0.89 | 0.79 |
| VECTOR | 0.13 | 0.59 | 0.57 | 0.96 | 0.98 | 0.98 |

Table 5: Benchmarking CVE-LLM against state-of-the-art Open Source LLMs. CVE-LLM = CVE-LLM results on test dataset (N=44K), CVE-LLM-Eval = deployed CVE-LLM results on test dataset (N=44K), CVE-LLM-Prod = post-deployment results on production(N=10K). Evaluation metrics are Rouge-L score for Internal and Customer comments, and micro-F1 score for the rest: for both metrics, higher numbers indicate better performance.

- If VEXJUSTIFICATION is ‘Other’ and Customer Comment is empty, then it is an error in generation. We use Internal Comment as External Comment as well in case it is not empty.
- If VEXCATEGORY is Affected, VEXJUSTIFICATION is automatically set to None, and it is ensured that the Vector has the proper CVSS format.

Model Result Integration Into CSMS System The model results are generated daily using vLLM architecture (Kwon et al. 2023) on one Tesla A100 SXM4 GPU via batch processing, using the newly arrived notifications from third party component vendors and software versions of SHS products. The AI evaluations are generated for every applicable product that uses the third party components relevant to the notifications, as well as for every past notification that is applicable to the product. These evaluations are stored in a database and made available to product cybersecurity experts who can auto-populate the CSMS with AI-recommended evaluations and/or revise the evaluations as needed. The corrected evaluations are subsequently archived for further processing for further training, in addition to investigating mitigation and determining the appropriate reply to customer.

Results

The experiment results are reported on two datasets: 1. Test dataset with 44K evaluations. 2. Post-deployment dataset results (CVE-LLM-Prod) with 10K expert evaluations.

We use ROUGE-L and micro-F1 for evaluating the model responses: ROUGE-L(Lin and Och 2004) is used for evaluating the model generated INTERNALCOMMENT and CUSTOMERCOMMENT, whereas micro-F1(Pedregosa et al. 2011) is used for VEXCATEGORY, VEXJUSTIFICATION and (environmental) Vector.

Benchmarking With Open Source LLMs

We have used open source models that are of similar size as CVE-LLM and are top-performing models on LLM Leaderboards (Huggingface 2023). We trained them using the DAPT and SFT methods described in the Model Training section. In addition, we also performed RAG over our test dataset with Llama3-70B. Table 5 enumerates the benchmarking results on our test dataset.

CVE-LLM, which is based on MPT-7B model, has shown better performance than Llama2-7B and Mistral-7B-based models. The highest performance improvement is observed for classification-based generations: VEXCATEGORY, VEXJUSTIFICATION and Vector. Post processing the model results have shown the highest improvement for VEXJUSTIFICATION and Vector. We also observe that the performance of CVE-LLM post-deployment (i.e., with completely new assets and notifications) is comparable to the test dataset results.

Ablation Studies

We conducted ablation studies with respect to different components of model training and inference: (a) DAPT-SFT vs SFT-only system, and (b) Impact of beam size, temperature, and nucleus sampling during model inference. For each of these experiments, the test dataset consists of the same data points from the Evaluations dataset. The training data points have varied across the datasets, but we have made sure that all the possible assets are represented in the training datasets.

Effect of Domain Adaptation The results after Supervised Finetuning of Domain-adapted CVE-LLM-Base vs Supervised Finetuning of MPT-7B is shown in Table 6. Domain adaptation and vocabulary expansion leads to a much better performance for all the instruction types, as shown in Table 6.

Effect of Inference Parameters In generative model inference, parameters such as temperature, beam size, and nucleus sampling significantly influence the quality and diversity of the generated outputs. Temperature controls the randomness of the output, with higher values leading to more diverse results. Nucleus Sampling (top-p sampling) dynamically adjusts the candidate set of tokens to balance between diversity and coherence based on the cumulative probability threshold ‘p’. Beam Size in beam search influences the depth of search for probable sequences, balancing between computational cost and output quality. Generally these parameters are adjusted to optimize the trade-off between diversity and coherence of the generated text. We performed inference on our test dataset using CVE-LLM-Eval model to conduct the experiment on the effect of inference parameters.

Temperature and Nucleus Sampling. In order to test the effect of temperature, we maintained the generation param-

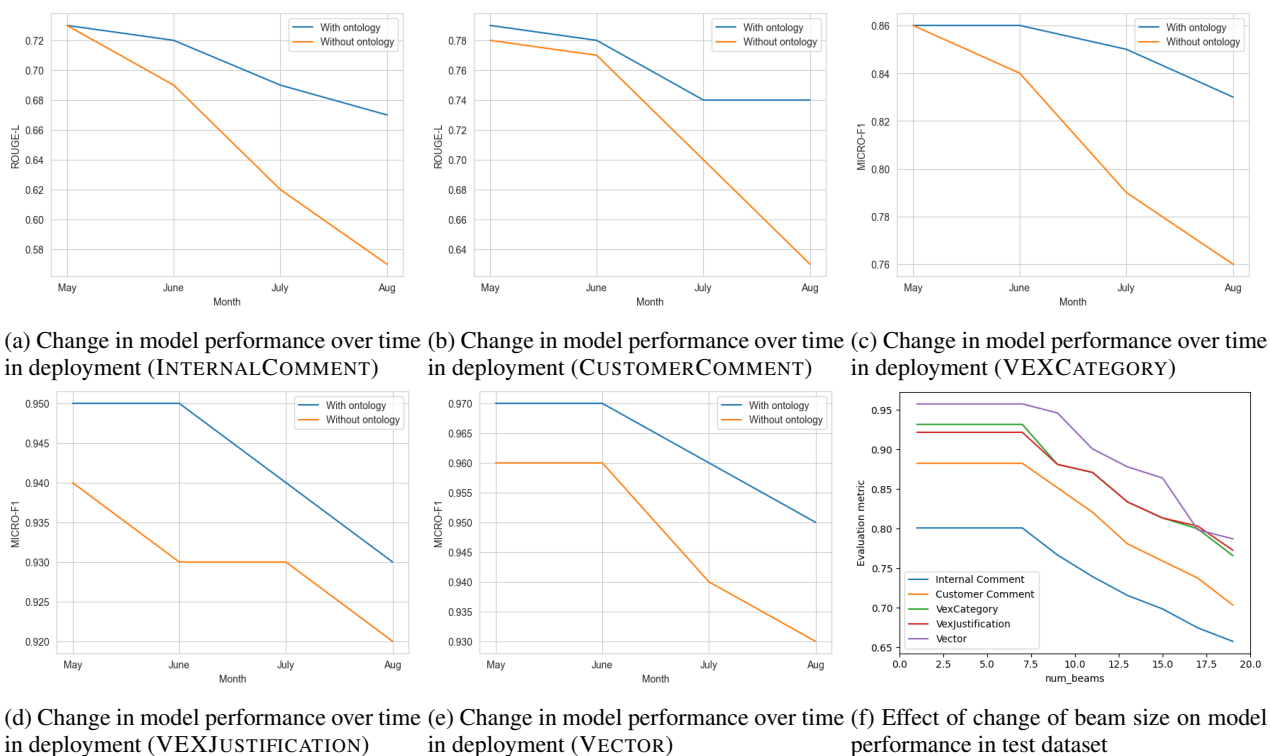


Figure 2: Ablation studies. The Evaluation metrics is Rouge-L for INTERNALCOMMENT and CUSTOMERCOMMENT, and micro-F1 for the rest.

eters at their default settings and conducted tests by varying the temperature values from 0 to 1 in increments of 0.1 to evaluate the outcomes. Similarly, for top-p sampling too, we changed the parameter ‘top-p’ from 0 to 1 in increments of 0.1, keeping all other generation parameters in default settings. For both of the experiments, no changes were observed.

Beam Size. We measured the effect of beam size for each of the instruction types by changing the beam size from 1 to 19 in increments of 2, maintaining other generation parameters in their default values. We notice that beyond beam size 7, the performance of the system steadily declines for all the instruction types. The effect of change of beam size on model generation is shown in Figure 2f.

Effect of Ontology-Assisted Enrichment. We have tested the effect of ontology enrichment on our test dataset (Table 6) as well as deployment data (Figure 2) with 10K evaluations. The model has been trained with data until April and has been deployed since May. The effect in deployment shows the efficacy of the ontology-assisted enrichment.

Inference Time

We have used different techniques to optimize for the time required in inference. The first experiment used popular model serving algorithm vLLM adapted for MPT-7B to find required inference time for one Tesla A100 SXM4 GPU for each of the instruction types and found speedup of 10x magnitude. The experiment results are shown in Table 7. In addition,

the adaptive generation by dividing the inference batch into batches of small and large token lengths also improves the average performance over the test dataset to 4.5 second per evaluation from 8 second. The other adaptation of using different maximum sequence length for each instruction type brings down the inference time for each data point in the Evaluations dataset by 1.5 times. We have calculated the time required by a cybersecurity expert in our organization to evaluate vulnerabilities. The mean time required for an expert, averaged over 7K evaluations is 194s with a median of 58s, whereas an evaluation by an our model takes only 1 second per evaluation.

| Evaluation | CVE-LLM | No DAPT | No ontology |
|------------------|---------|---------|-------------|
| VEXCATEGORY | 0.94 | 0.83 | 0.93 |
| VEXJUSTIFICATION | 0.90 | 0.80 | 0.89 |
| Internal Comment | 0.79 | 0.64 | 0.80 |
| Customer Comment | 0.88 | 0.69 | 0.88 |
| Vector | 0.96 | 0.84 | 0.95 |

Table 6: Test dataset Ablation Results: Results without DAPT and without ontology. The Evaluation metrics is Rouge-L for Internal and Customer comment, and micro-F1 for the rest.

| Evaluation | No speedup | vLLM |
|------------------|------------|------|
| VEXCATEGORY | 0.5 | 0.04 |
| Internal Comment | 2.5 | 0.20 |
| Customer Comment | 2 | 0.19 |
| Vector | 3 | 0.32 |

Table 7: Inference time in seconds on our test dataset for model serving with sequence length = 20K.

Discussion

In this section, we will elucidate our observations related to training and performance of LLMs for vulnerability evaluation, potential areas of improvement and future work.

Ontology Enrichment. While ontology enrichment did not yield substantial performance improvement on the test dataset, it demonstrated a significant impact during deployment on unseen notifications and assets. Most of the errors on unseen notifications pertain to the absence of a linking factor between a new notification and an older similar notification. A close examination show that for most notifications with similar prerequisites, the responses to the instructions are similar across assets. In the future, we would like to experiment with adding more knowledge about the notifications from other security knowledge graphs.

Domain Adaptation. We observed significant improvement in performance with domain adaptation of the LLM. Error analysis for the model without domain adaptation shows more hallucinations, more generalized comment generation and lower performance for VEXJUSTIFICATION and Vectors that are less represented.

Performance Across Instructions. The performance of CVE-LLM in classification tasks is generally better than the generation tasks. However, though the F1 score for Vector classification is generally high, it remains low for underrepresented classes, resulting in a decline in overall performance. Similarly, performance drops from VEXCATEGORY to VEXJUSTIFICATION classification, largely due to the presence of underrepresented classes in VEXJUSTIFICATION.

We surmise that the performance of CVE-LLM for CUSTOMERCOMMENT is better than that in INTERNALCOMMENT due to the following factors:

- We observed that approximately 85% of the customer comments follow a highly templated format and are repeated across multiple assets multiple subdivisions for similar notifications, which is not the case for internal comments.
- Internal comments tend to mention specific products, components and versions, which is generally not seen in customer comments. CVE-LLM tends to hallucinate on named entities.

Generation Parameters. We have noticed model output variations with beam size, with decline in performance with higher beam sizes. Increase in beam search size increases the

diversity of text, and that leads to less contextually relevant outputs, as seen in our experiments. Low beam size is sufficient for high quality generation because of the specialized nature of the dataset. The difference between probabilities of less probable outcomes and highly probable outcomes, given the context, is low. This leads to no temperature or nucleus sampling-based variations in model-generated outputs.

Error Analysis Across LLMs

We have used MPT-7B to train CVE-LLM for two reasons: (a) Expansion capability to at least 20K tokens without significant loss of performance, and (b) It performs better than other open source models we tested with our dataset. However, among all the LLMs we have trained, majority of the issues we have encountered are remarkably similar, differing primarily in their severity. The errors we have encountered across LLMs can be of four types:

- **Hallucination in named entities:** The generated text for both Internal Comment and Customer Comment has been observed to have omitted or falsely included several critical details, particularly those concerning the affected software versions, the recommended update versions, the names of the software or components, and other similarly named entities that constitute essential information. We plan to introduce custom losses during model training in the future that will counteract these hallucinations.
- **Spurious text generation:** This error is more common for Llama2 model and was observed in MPT-7B before introduction of the dedicated STOP token.
- **Performance degradation due to long text:** In spite of having the ability to extend the sequence length of the trained model beyond the training sequence length, the performance on long sequences still falls behind. Long instructions comprises of multiple vulnerabilities, and constitute only 10% of our test dataset. In the future, we plan to segment it into multiple vulnerabilities and use more robust semantic understanding of each vulnerability to perform vulnerability chaining. We also tested the system only with zero shot instructions at the inference, and we surmise that we can utilize language understanding capabilities better by using few shot methods to assess vulnerabilities for assets.
- **Errors due to different patterns of task responses in organization subdivisions:** The variation of task responses across subdivisions have been seen to affect the quality of generation though the instruction input context includes the subdivision name. This is an issue especially when a notification has vastly more representation in another subdivision. In the future we would attempt to address this issue with data augmentation and sampling techniques.

Relevant Work

Language Models have been used extensively in vulnerability management (a) to determine CVSS metrics from CVE description, (b) to establish a mapping between vulnerabilities in CVE database with the corresponding attack tac-

tics and techniques in ATT&CK database ⁶, (c) for vulnerability detection, and (d) for vulnerability repair. Determination of CVSS metrics using vulnerability description has been treated as a text classification/regression problem. Mapping to CVSS score follows a linear regression using Bag-of-Words model (Elbaz, Rilling, and Morin 2020) or neural network model using Doc2Vec (Vasireddy, Dale, and Li 2023) method of extracting features from a document. CVSS-BERT (Shahid and Debar 2021), on the other hand, trained different BERT models for classification of CVE descriptions to the values of different CVE vectors. Both encoder and decoder based models have been used for mapping vulnerability to ATT&CK tactics and techniques. CVET (Ampel et al. 2021), a RoBERTa(Liu et al. 2019)-based model, classified CVE descriptions to one of the ten tactics in ATT&CK, whereas SMET (Abdeen et al. 2023) used BERT-based textual similarity to map CVE entries to ATT&CK techniques. Though ChatGPT-based approaches did not yield state-of-the-art results (Liu et al. 2023), VTT-LLM (Zhang et al. 2024), trained using various versions of the decoder Bloom model(Le Scao et al. 2023) with Chain-of-thought (Wei et al. 2022) instructions, incorporated relations between core concepts in CWE (Christey et al. 2013) and CAPEC databases ⁷ for CVE to ATT&CK mapping and surpassed encoder based models. Other work (Yosifova, Tashveva, and Trifonov 2021) involved CVE Vulnerability type classification using TF-IDF and standard machine learning classifier.

For vulnerability detection, encoder-based LLMs (Ameri et al. 2021; Yin et al. 2020) have been instrumental, particularly through innovative approaches like the pretrain-and-finetune paradigm. These approaches have also leveraged novel pretraining strategies, the integration Graph Neural Networks(Sewak, Emani, and Naresh 2023) or Long Short-Term Memory (LSTM)(Hassanin et al. 2024) networks, prompt tuning, program analysis, specialized calculus-based causal reasoning, and knowledge graph-based reasoning. In recent times, the focus has shifted towards decoder-only LLMs(Zhou et al. 2024) for vulnerability detection. Our work focuses on using pretrain-and-finetune approach with decoder-based LLMs, coupled with ontology enrichment for vulnerability evaluation generation.

Conclusion

This study illustrates the capacity of large language models (LLMs) to learn from expert-curated historical vulnerability evaluation data, thereby enabling the automation of vulnerability evaluation generation for MDMs. Our model (CVE-LLM) has proven to be effective in learning from past evaluations and has shown high accuracy in prediction of VEXCATEGORY and CVSS Vectors. The model's inference is also substantially faster, achieving speeds approximately 50-100 times greater than those of a human expert. The system is deployed in a human-in-the-loop manner, assisting product cybersecurity experts in swiftly identifying vulnerabilities that affect the assets and communicating mitigations

⁶MITRE ATT&CK database: <https://attack.mitre.org/>

⁷MITRE CAPEC database: <https://capec.mitre.org/>

promptly to the customer.

In future work, we intend to investigate additional knowledge infusion techniques, incorporating a broader range of cybersecurity databases. Addressing the challenge of hallucinations remains a key focus. Furthermore, we aim to develop mechanisms for integrating more comprehensive product knowledge by leveraging source code and official cybersecurity documentation. Beyond standardized databases, numerous blogs and websites provide continuous updates on vulnerabilities, which we intend to incorporate into our DAPT and SFT training paradigms. The success of CVE-LLM in the SHS vulnerability evaluation platform has also highlighted opportunities to assist cybersecurity experts in vulnerability mitigation.

Disclaimer

The concepts and information presented in this paper/presentation are based on research results that are not commercially available. Future commercial availability cannot be guaranteed.

References

- Abdeen, B.; Al-Shaer, E.; Singhal, A.; Khan, L.; and Hamlen, K. 2023. Smet: Semantic mapping of cve to att&ck and its application to cybersecurity. In *IFIP Annual Conference on Data and Applications Security and Privacy*, 243–260. Springer.
- Ameri, K.; Hempel, M.; Sharif, H.; Lopez Jr, J.; and Perumalla, K. 2021. Cybert: Cybersecurity claim classification by fine-tuning the bert language model. *Journal of Cybersecurity and Privacy*, 1(4): 615–637.
- Ampel, B.; Samtani, S.; Ullman, S.; and Chen, H. 2021. Linking common vulnerabilities and exposures to the mitre att&ck framework: A self-distillation approach. *arXiv preprint arXiv:2108.01696*.
- Chen, X.; Liang, C.; Huang, D.; Real, E.; Wang, K.; Pham, H.; Dong, X.; Luong, T.; Hsieh, C.-J.; Lu, Y.; et al. 2024. Symbolic discovery of optimization algorithms. *Advances in Neural Information Processing Systems*, 36.
- Christey, S.; Kenderdine, J.; Mazella, J.; and Miles, B. 2013. Common weakness enumeration. *Mitre Corporation*.
- Deng, G.; Liu, Y.; Mayoral-Vilches, V.; Liu, P.; Li, Y.; Xu, Y.; Zhang, T.; Liu, Y.; Pinzger, M.; and Rass, S. 2024. PentestGPT: Evaluating and Harnessing Large Language Models for Automated Penetration Testing. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Elbaz, C.; Rilling, L.; and Morin, C. 2020. Fighting N-day vulnerabilities with automated CVSS vector prediction at disclosure. In *Proceedings of the 15th International Conference on Availability, Reliability and Security*, 1–10.
- Happe, A.; and Cito, J. 2023. Getting pwn'd by ai: Penetration testing with large language models. In *Proceedings of*

- the 31st ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering, 2082–2086.
- Hassanin, M.; Keshk, M.; Salim, S.; Alsubaie, M.; and Sharma, D. 2024. PLLM-CS: Pre-trained Large Language Model (LLM) for Cyber Threat Detection in Satellite Networks. *arXiv preprint arXiv:2405.05469*.
- Huggingface. 2023. Open LLM Leaderboard.
- Iannacone, M.; Bohn, S.; Nakamura, G.; Gerth, J.; Huffer, K.; Bridges, R.; Ferragut, E.; and Goodall, J. 2015. Developing an ontology for cyber security knowledge graphs. In *Proceedings of the 10th annual cyber and information security research conference*, 1–4.
- Kiesling, E.; Ekelhart, A.; Kurniawan, K.; and Ekaputra, F. 2019. The SEPSES knowledge graph: an integrated resource for cybersecurity. In *International Semantic Web Conference*, 198–214. Springer.
- Kwon, W.; Li, Z.; Zhuang, S.; Sheng, Y.; Zheng, L.; Yu, C. H.; Gonzalez, J.; Zhang, H.; and Stoica, I. 2023. Efficient memory management for large language model serving with pagdatention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, 611–626.
- Le Scao, T.; Fan, A.; Akiki, C.; Pavlick, E.; Ilić, S.; Hesslow, D.; Castagné, R.; Luccioni, A. S.; Yvon, F.; Gallé, M.; et al. 2023. Bloom: A 176b-parameter open-access multilingual language model.
- Li, Y.; Wang, S.; Ding, H.; and Chen, H. 2023. Large language models in finance: A survey. In *Proceedings of the Fourth ACM International Conference on AI in Finance*, 374–382.
- Lin, C.-Y.; and Och, F. 2004. Looking for a few good metrics: ROUGE and its evaluation. In *Ntcir workshop*.
- Liu, X.; Tan, Y.; Xiao, Z.; Zhuge, J.; and Zhou, R. 2023. Not the end of story: An evaluation of chatgpt-driven vulnerability description mappings. In *Findings of the Association for Computational Linguistics: ACL 2023*, 3724–3731.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Mell, P.; Scarfone, K.; and Romanosky, S. 2006. Common vulnerability scoring system. *IEEE Security & Privacy*, 4(6): 85–89.
- Mitre, C. 2005. Common vulnerabilities and exposures.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, E. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research*, 12: 2825–2830.
- Rasley, J.; Rajbhandari, S.; Ruwase, O.; and He, Y. 2020. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 3505–3506.
- Sallam, M. 2023. ChatGPT utility in healthcare education, research, and practice: systematic review on the promising perspectives and valid concerns. In *Healthcare*, volume 11, 887. MDPI.
- Sewak, M.; Emani, V.; and Naresh, A. 2023. CRUSH: Cybersecurity Research using Universal LLMs and Semantic Hypernetworks.
- Shahid, M. R.; and Debar, H. 2021. Cvss-bert: Explainable natural language processing to determine the severity of a computer security vulnerability from its description. In *2021 20th IEEE International Conference on Machine Learning and Applications (ICMLA)*, 1600–1607. IEEE.
- Syed, Z.; Padia, A.; Finin, T.; Mathews, L.; and Joshi, A. 2016. UCO: A unified cybersecurity ontology. In *Workshops at the thirtieth AAAI conference on artificial intelligence*.
- Taori, R.; Gulrajani, I.; Zhang, T.; Dubois, Y.; Li, X.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*. <https://crfm.stanford.edu/2023/03/13/alpaca.html>, 3(6): 7.
- Team, M. N. 2023. Introducing MPT-7B: A New Standard for Open-Source, Commercially Usable LLMs. Accessed: 2023-05-05.
- Vasireddy, D. T.; Dale, D. S.; and Li, Q. 2023. CVSS Base Score Prediction Using an Optimized Machine Learning Scheme. In *2023 Resilience Week (RWS)*, 1–6. IEEE.
- Wei, J.; Wang, X.; Schuurmans, D.; Bosma, M.; Xia, F.; Chi, E.; Le, Q. V.; Zhou, D.; et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35: 24824–24837.
- Wu, Q.; Bansal, G.; Zhang, J.; Wu, Y.; Zhang, S.; Zhu, E.; Li, B.; Jiang, L.; Zhang, X.; and Wang, C. 2023. Autogen: Enabling next-gen llm applications via multi-agent conversation framework. *arXiv preprint arXiv:2308.08155*.
- Yang, X.; Wang, Z.; Wang, Q.; Wei, K.; Zhang, K.; and Shi, J. 2024. Large language models for automated q&a involving legal documents: a survey on algorithms, frameworks and applications. *International Journal of Web Information Systems*.
- Yin, J.; Tang, M.; Cao, J.; and Wang, H. 2020. Apply transfer learning to cybersecurity: Predicting exploitability of vulnerabilities by description. *Knowledge-Based Systems*, 210: 106529.
- Yosifova, V.; Tasheva, A.; and Trifonov, R. 2021. Predicting vulnerability type in common vulnerabilities and exposures (cve) database with machine learning classifiers. In *2021 12th National Conference with International Participation (ELECTRONICA)*, 1–6. IEEE.
- Zhang, C.; Wang, L.; Fan, D.; Zhu, J.; Zhou, T.; Zeng, L.; and Li, Z. 2024. VTT-LLM: Advancing Vulnerability-to-Tactic-and-Technique Mapping through Fine-Tuning of Large Language Model. *Mathematics*, 12(9): 1286.
- Zhou, X.; Cao, S.; Sun, X.; and Lo, D. 2024. Large Language Model for Vulnerability Detection and Repair: Literature Review and Roadmap. *arXiv preprint arXiv:2404.02525*.