

# Early-Stage Identification of Exploit-Prone Vulnerabilities Using Parameterized Machine Learning Models

**<sup>1</sup>Deepanshu Sharma,<sup>2</sup>Dr. Inderpal Singh Oberoi**

<sup>1</sup>Research Scholar, Department of Computer Applications, Maharaja Agrasen Himalayan Garhwal University

<sup>2</sup>Assistant Professor, Department of Computer Applications, Maharaja Agrasen Himalayan Garhwal University

## ABSTRACT

The rapid growth of software systems has been accompanied by a steady increase in reported security vulnerabilities, creating significant challenges for organizations attempting to prioritize mitigation efforts. Since only a small subset of disclosed vulnerabilities are eventually exploited, early identification of exploit-prone vulnerability management is critical for effective vulnerability management. This study proposes a parameterized machine learning approach for predicting exploit likelihood at the time of vulnerability disclosure, relying exclusively on static vulnerability parameters available at early stages. Using features derived from Common Vulnerability Scoring System (CVSS) metrics and disclosure metadata, multiple supervised classification models are developed and evaluated. The results demonstrate that parameterized machine learning models can achieve meaningful predictive accuracy without relying on post-disclosure or exploit-availability data. The findings highlight the practical value of early-stage exploit prediction for secure software development, proactive defense, and efficient patch prioritization.

**Keywords:** Exploit Prediction; Software Vulnerabilities; Machine Learning; CVSS; Early-Stage Security Assessment; Patch Prioritization; Secure Software Development

## 1. INTRODUCTION

Software vulnerabilities remain one of the most critical threats to information security. Each year, thousands of vulnerabilities are disclosed through public repositories such as the National Vulnerability Database (NVD). However, empirical evidence consistently shows that only a fraction of disclosed vulnerabilities are ever exploited in real-world attacks, while security teams are forced to respond to all disclosures with limited resources. This imbalance creates an urgent need for accurate methods to identify exploit-prone vulnerabilities as early as possible.

Traditional vulnerability management practices rely heavily on severity scores such as CVSS or on the later appearance of public exploits. While severity scores provide a general assessment of potential impact, they are not designed to predict attacker behavior. Similarly, waiting for exploit code to appear defeats the purpose of early risk mitigation. As a result, organizations often misallocate patching resources, focusing on highly scored but

rarely exploited vulnerabilities while overlooking lower-scored vulnerabilities that are actively exploited.

Machine learning has emerged as a promising approach for vulnerability exploit prediction. Most existing studies, however, rely on temporal or dynamic features, such as exploit availability, social media activity, or post-disclosure trends. These features are not available at the time of vulnerability disclosure, limiting their usefulness for early-stage decision-making.

This paper addresses this gap by proposing a parameterized machine learning framework that predicts exploit likelihood at disclosure between CVSS severity and real-world exploitation.

Machine learning-based approaches have demonstrated improved performance by learning patterns from historical vulnerability data. Sabottke et al. showed that social media signals could improve exploit prediction, while other studies incorporated exploit database references or temporal trends. Although these approaches yield high accuracy, they depend on information that becomes available only after disclosure.

A smaller body of work has explored early-stage exploit prediction, focusing on static features such as vulnerability type, attack vector, and impact metrics. These studies suggest that attacker preferences can be partially inferred from intrinsic vulnerability characteristics. However, comprehensive evaluations of parameterized models using only static features remain limited.

This research contributes to the literature by systematically analyzing how static vulnerability parameters, when combined

time using only static vulnerability parameters. The study evaluates whether early-available features can provide sufficient predictive power to support secure software development and effective patch prioritization.

## 2. RELATED WORK

Prior research on vulnerability exploit prediction can be broadly categorized into severity-based, temporal-based, and machine learning-based approaches. Severity-based methods rely primarily on CVSS scores, assuming that higher severity implies higher exploit likelihood. However, multiple studies have shown weak correlation between CVSS severity and exploit likelihood. Machine learning-based approaches, with parameterized machine learning models, can support reliable exploit prediction at disclosure time.

## 3. RESEARCH OBJECTIVES

The main objectives of this study are:

1. To develop machine learning models capable of predicting exploit likelihood at the time of vulnerability disclosure.
2. To evaluate the effectiveness of static vulnerability parameters in exploit prediction.
3. To compare the performance of different classifier algorithms under a parameterized modeling framework.
4. To analyze the practical implications of early-stage exploit prediction for secure software development and patch prioritization.

## 4. METHODOLOGY

### 4.1 Dataset Description

The study uses publicly available vulnerability data collected from the National Vulnerability Database (NVD).

Exploit labels are assigned by matching CVE identifiers with known exploit references from exploit repositories. Vulnerabilities are classified into two categories: exploit-prone and non-exploit-prone.

#### 4.2 Feature Selection: Static Vulnerability Parameters

Only features available at disclosure time are considered. These include:

- **CVSS Base Metrics:** Attack Vector, Attack Complexity, Privileges Required, User Interaction, Confidentiality Impact, Integrity Impact, Availability Impact, and Base Score.
- **Vulnerability Metadata:** Vulnerability type, affected platform, and disclosure year.
- **Structural Parameters:** Categorical encoding of vulnerability class (e.g., buffer overflow, injection, access control).

No temporal, exploit-availability, or post-disclosure features are used.

#### 4.3 Machine Learning Models

The following supervised classifiers are implemented:

- Logistic Regression
- Support Vector Machine (SVM)
- Random Forest
- Gradient Boosting (XGBoost)

Each model is parameterized using optimized hyperparameters determined through cross-validated search techniques.

#### 4.4 Evaluation Metrics

Model performance is evaluated using:

- Accuracy
- Precision
- Recall

- F1-Score

- Area Under the ROC Curve (AUC)

Stratified k-fold cross-validation is used to ensure robustness.

### 5. RESULTS AND ANALYSIS

The experimental results indicate that parameterized machine learning models can effectively identify exploit-prone vulnerabilities at early stages.

Random Forest and Gradient Boosting models outperform linear classifiers, achieving higher F1-scores and AUC values. The inclusion of detailed CVSS base metrics significantly improves model performance compared to using overall severity scores alone. Results further show that attack vector and attack complexity are among the most influential predictors, reflecting attacker preference for remotely exploitable and low-complexity vulnerabilities.

Despite relying exclusively on static parameters, the best-performing models demonstrate strong discriminative ability, confirming the feasibility of early-stage exploit prediction.

### 6. DISCUSSION

The findings of this study highlight the importance of **feature granularity and model parameterization** in vulnerability exploit prediction. While static features do not capture attacker behavior directly, they encode meaningful signals about exploit feasibility and potential reward. Machine learning models are able to learn these signals and generalize across vulnerability classes.

From a practical perspective, early-stage exploit prediction can significantly enhance secure software development lifecycles. Developers can prioritize code reviews and

testing efforts for vulnerabilities predicted to be exploit-prone, while security teams can allocate patching resources more effectively. The study also demonstrates that meaningful security intelligence can be extracted without relying on external or post-disclosure data, making the approach suitable for real-time vulnerability assessment systems.

## 7. CONCLUSION

This research has demonstrated that the early-stage identification of exploit-prone software vulnerabilities is both feasible and effective when using parameterized machine learning models that rely exclusively on static vulnerability parameters. Contrary to conventional practices that depend on post-disclosure indicators such as exploit availability, attack trends, or temporal signals, this study confirms that meaningful predictive intelligence can be extracted at the moment of vulnerability disclosure. By utilizing structured information embedded within CVSS base metrics and disclosure-time metadata, machine learning models can infer exploit likelihood with a high degree of reliability.

The findings clearly show that granular CVSS base metrics are substantially more informative than aggregate severity scores alone. Parameters such as attack vector, attack complexity, required privileges, and user interaction play a decisive role in determining exploit feasibility. These parameters effectively encode attacker effort, accessibility, and potential impact—key factors influencing real-world exploitation decisions. When processed through parameterized learning models, these static attributes collectively enable

accurate classification of exploit-prone vulnerabilities, even in the absence of exploit code or attacker activity data.

A key contribution of this research lies in its demonstration that early-stage exploit prediction supports proactive security decision-making. Security teams are often constrained by limited resources and must prioritize mitigation efforts across thousands of disclosed vulnerabilities. Early predictive insights allow organizations to focus on vulnerabilities with a higher probability of exploitation, thereby improving the efficiency of patch management processes. This approach reduces the risk window between vulnerability disclosure and remediation, which is critical in preventing large-scale security breaches.

From a software engineering perspective, the results highlight the importance of integrating exploit prediction mechanisms into the secure software development lifecycle (SSDLC). Early identification of exploit-prone weaknesses enables developers to prioritize secure coding practices, targeted code reviews, and rigorous testing for high-risk components. This proactive stance enhances overall software resilience and reduces long-term maintenance and incident response costs.

The research also demonstrates the practical applicability of parameterized machine learning models in real-world environments. Because the proposed approach relies only on static features available at disclosure time, it is well-suited for automation within vulnerability assessment tools, security dashboards, and continuous integration pipelines. Organizations can deploy such models without relying on external threat

intelligence feeds or waiting for exploit confirmation, making early-stage risk assessment both scalable and cost-effective. Despite these promising results, the study acknowledges certain limitations. Static vulnerability parameters, while informative, cannot capture evolving attacker behavior, exploit chaining, or contextual deployment factors such as asset criticality and exposure. As such, early-stage predictions should be viewed as decision-support mechanisms rather than definitive indicators of exploitation. Nonetheless, their value in guiding prioritization and risk mitigation is substantial.

Future research directions include the integration of explainable artificial intelligence (XAI) techniques to improve transparency and trust in machine learning-based exploit prediction systems. Explainable models would allow security analysts to understand why certain vulnerabilities are classified as exploit-prone, facilitating better human-machine collaboration. Additionally, evaluating model performance across diverse software ecosystems, programming languages, and industrial domains would further strengthen generalizability. Combining static parameter-based models with controlled dynamic signals may also enhance prediction accuracy while preserving early-stage applicability.

This research establishes that parameterized machine learning models using static vulnerability parameters provide a viable, practical, and effective solution for early-stage exploit-prone vulnerability identification. By enabling timely and informed security decisions, this approach

contributes significantly to improved vulnerability management, stronger secure software development practices, and enhanced organizational cyber resilience.

## REFERENCES

1. Allodi, L. and Massacci, F. (2014) 'Comparing vulnerability severity and exploitability using CVSS', *IEEE Security & Privacy*, 12(1), pp. 52–60.
2. Arora, A., Telang, R. and Xu, H. (2008) 'Optimal policy for software vulnerability disclosure', *Management Science*, 54(4), pp. 642–656.
3. Bishop, C.M. (2006) *Pattern Recognition and Machine Learning*. Springer.
4. Breiman, L. (2001) 'Random forests', *Machine Learning*, 45(1), pp. 5–32.
5. Chowdhury, I. and Zulkernine, M. (2011) 'Using complexity metrics to predict software vulnerabilities', *Journal of Systems Architecture*, 57(3), pp. 294–313.
6. Feurer, M. and Hutter, F. (2019) 'Hyperparameter optimization', *Springer Series on Machine Learning*.
7. Fenton, N. and Neil, M. (1999) 'A critique of software defect prediction models', *IEEE Transactions on Software Engineering*, 25(5), pp. 675–689.
8. Friedman, J.H. (2001) 'Greedy function approximation: A gradient boosting machine', *Annals of Statistics*, 29(5), pp. 1189–1232.

9. Hastie, T., Tibshirani, R. and Friedman, J. (2009) *The Elements of Statistical Learning*. Springer.
10. Houmb, S.H., Franqueira, V.N.L. and Engum, E.A. (2010) 'Estimating software security risk', *Information and Software Technology*, 52(6), pp. 589–599.
11. Joachims, T. (1998) 'Text categorization with support vector machines', *ECML*, pp. 137–142.
12. Khoshgoftaar, T.M. and Allen, E.B. (2003) 'Logistic regression modeling of software quality', *IJRQSE*, 10(4), pp. 435–448.
13. Li, Y., Tan, K.L. and Li, Z. (2016) 'Predicting vulnerability exploitability using machine learning', *IEEE Software*, 33(5), pp. 43–51.
14. Mell, P., Scarfone, K. and Romanosky, S. (2007) 'A complete guide to the CVSS', *FIRST*.
15. Neuhaus, S. and Zimmermann, T. (2010) 'Security trend analysis with CVE topic models', *IEEE S&P*, pp. 111–125.
16. Ozment, A. (2007) 'Improving vulnerability discovery models', *ACM CCS*, pp. 327–338.
17. Provost, F. and Fawcett, T. (2013) *Data Science for Business*. O'Reilly.
18. Rescorla, E. (2005) 'Is finding security holes a good idea?', *IEEE Security & Privacy*, 3(1), pp. 14–19.
19. Sabottke, C., Suciu, O. and Dumitraş, T. (2015) 'Vulnerability disclosure in the age of social media', *USENIX Security*, pp. 1041–1056.
20. Scikit-learn Developers (2011) 'Scikit-learn: machine learning in Python', *JMLR*, 12, pp. 2825–2830.
21. Shin, Y. et al. (2011) 'Evaluating complexity, churn, and developer activity metrics', *IEEE TSE*, 37(6), pp. 772–787.
22. Sommer, R. and Paxson, V. (2010) 'Outside the closed world', *IEEE S&P*, pp. 305–316.
23. Sutton, R.S. and Barto, A.G. (1998) *Reinforcement Learning*. MIT Press.
24. Tsipenyuk, K., Chess, B. and McGraw, G. (2005) 'Seven pernicious kingdoms', *IEEE Security & Privacy*, 3(6), pp. 81–84.
25. Verendel, V. (2009) 'Quantified security is a weak hypothesis', *NSPW*, pp. 37–49.
26. Williams, L. and Wierman, M. (2010) 'Security in agile software development', *IEEE Software*, 27(3), pp. 14–16.
27. Zhang, H. et al. (2011) 'Measuring software security defects using complexity metrics', *Journal of Systems and Software*, 84(9), pp. 1608–1620.
28. Zimmermann, T. et al. (2010) 'Predicting defects using network analysis', *ICSE*, pp. 531–540.
29. Zou, C.C., Gong, W. and Towsley, D. (2002) 'Code red worm propagation modeling', *ACM CCS*, pp. 138–147.
30. Zulkernine, M. et al. (2010) 'Predicting vulnerabilities using software complexity metrics', *QSIC*, IEEE, pp. 23–32.