UNIOFFA JOURNAL OF INTERDISCIPLINARY STUDIES (U-OJIS)
A Multidisciplinary Journal of Science, Management and Social Sciences
Website: www.unioffajis.com
ISSN:3141-5571
E-ISSN:3141-558X

Predictive Vulnerability Management Using Artificial Intelligence for Critical Infrastructure in Cybersecurity

Jesse Egbhakhagbeobo Alegieuno
Security engineer, AI secure.
Student, department of information technology and management, school Illinois institute of technology
Jalegieuno@hawk.Illinoistech.edu

Dr. Maurice Dawson
CEO, AI secure consulting group,
Associate Professor, Illinois Institute of Illinois Institute of Technology
mdawson@aisecgrp.co ; mdawson2@illinoistech.edu

Volume 1:1

Abstract

The rapid growth of software vulnerabilities has created a major challenge for organisations responsible for protecting critical infrastructure because conventional vulnerability management approaches frequently generate remediation queues that exceed available technical and human resources. This study developed and empirically evaluated an Artificial Intelligence (AI)-based predictive vulnerability management framework for prioritising vulnerabilities according to their probability of exploitation, asset criticality, exposure, and potential attack-path relevance. A controlled simulation-based experimental design was adopted using a synthetic dataset of 20,000 vulnerability records containing vulnerability severity, Exploit Prediction Scoring System (EPSS) score, Known Exploited Vulnerability (KEV) status, public exploit availability, asset criticality, internet exposure, attack-surface characteristics, asset age, patch complexity, dependency count, attacker interest, and mitigation strength. Four machine learning algorithms, namely Random Forest, Support Vector Machine, XGBoost, and LightGBM, were trained and evaluated using accuracy, precision, recall, F1-score, ROC-AUC, and false-positive rate. The results showed that LightGBM achieved the highest predictive accuracy (78.50%) and ROC-AUC (0.8090), while Random Forest produced the highest recall (94.35%). Comparative baseline analysis showed that the AI models provided stronger balanced predictive performance than CVSS-only, EPSS-threshold, and KEV-only approaches. The study further developed a context-aware prioritisation mechanism that combines predicted exploitation likelihood with asset criticality, internet exposure, and confirmed exploitation evidence. The findings demonstrate that AI can support a transition from severity-based vulnerability management to predictive and context-aware remediation. The proposed framework provides a practical basis for identifying vulnerabilities most likely to be exploited, prioritising critical systems, and informing attack-path-oriented defensive decisions within critical infrastructure environments.

Keywords: Artificial Intelligence, Vulnerability Management, Machine Learning, Cybersecurity, Critical Infrastructure, CVE, EPSS, CVSS, MITRE ATT&CK, Predictive Analytics.

1. Introduction

Critical infrastructure organisations increasingly depend on interconnected digital systems to deliver essential services. Energy, telecommunications, transportation, healthcare, financial services, water systems, government services, and industrial control environments rely on complex combinations of information technology and operational technology. This digital interdependence has simultaneously increased the potential consequences of cybersecurity incidents and expanded the number of vulnerabilities that security teams must manage. The central difficulty is no longer simply identifying vulnerabilities. Modern organisations must decide which vulnerabilities deserve immediate remediation when hundreds or thousands remain unresolved. CISA has specifically recognised that

ISSN:3141-5571
E-ISSN:3141-558X

organisations face difficulties determining which vulnerabilities pose the greatest practical risk and established the Known Exploited Vulnerabilities catalogue to focus attention on vulnerabilities for which exploitation has been observed.

The Common Vulnerability Scoring System (CVSS) remains an important mechanism for communicating technical vulnerability severity. However, severity does not necessarily indicate whether attackers are likely to exploit a vulnerability. Research on vulnerability remediation has demonstrated that machine learning can improve prioritisation by incorporating multiple characteristics associated with exploitation in the wild rather than relying on a single severity score (Jacobs et al., 2021; Peng et al., 2020).

The development of the Exploit Prediction Scoring System (EPSS) represents an important advancement because it estimates the probability that a vulnerability will be exploited. EPSS is now a data-driven machine-learning system that incorporates extensive signals from vulnerability and threat environments and produces exploitation probabilities for published CVEs.

Nevertheless, exploitation probability alone is insufficient for critical infrastructure. A vulnerability with moderate exploitation probability on a non-critical internal workstation may present less organisational risk than a similarly scored vulnerability affecting an internet-facing energy-management server. Effective vulnerability management therefore requires integration of vulnerability characteristics, exploitation likelihood, asset criticality, exposure, and potential attack paths.

MITRE ATT&CK provides an appropriate framework for incorporating adversarial behaviour because it maps tactics and techniques used by threat actors across the attack lifecycle. MITRE's Critical Infrastructure Cyberspace Analysis Tool further demonstrates how infrastructure models, ATT&CK techniques, CVEs and attack-path analysis can be combined to identify plausible routes through critical infrastructure environments. Against this background, this study develops an AI-based predictive vulnerability management framework designed to answer three questions: which vulnerabilities are most likely to be exploited, which systems should receive remediation priority, and how vulnerability information can be incorporated into attack-path analysis. The study contributes an empirical machine-learning comparison and a context-aware prioritisation framework for critical infrastructure cybersecurity.

2. Literature Review

Conceptual Review

Vulnerability Management in Cybersecurity

Vulnerability management is the continuous process through which organisations identify, assess, prioritise, remediate, and monitor security weaknesses in information systems. Traditional approaches generally begin with vulnerability scanning followed by ranking according to technical severity. CVSS has become an important standard for expressing vulnerability severity, but severity-based prioritisation can produce large numbers of high-risk findings without adequately distinguishing vulnerabilities that attackers are actually likely to exploit.

This limitation has encouraged a shift from vulnerability assessment towards vulnerability prioritisation. Peng et al. (2020) demonstrated that machine-learning approaches can improve remediation strategies by incorporating multiple vulnerability characteristics and exploitation evidence. Similarly, context-aware vulnerability prioritisation research argues that vulnerability risk changes according to temporal and organisational circumstances and should therefore incorporate contextual information rather than treating CVEs as isolated technical objects (Hemberg et al., 2022).

Artificial Intelligence and Predictive Vulnerability Management

AI changes vulnerability management by allowing systems to learn relationships between vulnerability characteristics and subsequent exploitation. Instead of assigning priority exclusively through predefined rules, machine-learning models can estimate the probability that a vulnerability will become an actual security incident.

Supervised algorithms such as Random Forest, Support Vector Machines, gradient boosting and neural networks are particularly appropriate where historical exploitation data are available. The EPSS framework provides a prominent real-world example of this approach. Its original research reported that machine learning could estimate exploitation probability and provide a more empirically grounded prioritisation mechanism than subjective severity assessment alone (Jacobs et al., 2021).

ISSN:3141-5571
E-ISSN:3141-558X

CVE, CVSS, EPSS and KEV

The Common Vulnerabilities and Exposures (CVE) system provides identifiers for publicly disclosed vulnerabilities, while the National Vulnerability Database enriches vulnerability information with technical metadata. CISA's KEV catalogue adds another important dimension by identifying vulnerabilities for which exploitation has been observed.

EPSS complements these resources by estimating future exploitation probability. Unlike a static severity score, EPSS is continuously updated and incorporates empirical signals associated with attacker behaviour. The combination of CVSS, EPSS and KEV is therefore potentially more informative than reliance on any single measure.

Recent research has explicitly explored this combination. A 2026 IEEE Access study proposed a vulnerability-management chaining framework combining KEV, EPSS and CVSS to reduce remediation workload while retaining coverage of exploited vulnerabilities.

Vulnerability Prioritisation and Critical Infrastructure

Critical infrastructure requires contextual vulnerability prioritisation because the consequences of compromise differ considerably across assets. Asset criticality, internet exposure, operational dependency, network position and potential effect on essential services should therefore influence remediation decisions.

CISA guidance recommends prioritising known exploited vulnerabilities on internet-facing systems and applying compensating controls where patching could compromise operational availability or safety. This is particularly important in operational technology environments where immediate patching may itself introduce safety and availability risks.

Attack-Path Analysis

A vulnerability may be relatively insignificant when considered independently but highly dangerous when it forms part of a chain leading to a critical asset. Attack-path analysis addresses this problem by examining how an adversary could move from an initial access point through vulnerable systems towards a valuable target.

MITRE's CICAT demonstrates this principle by combining infrastructure models with ATT&CK techniques and CVEs to generate plausible attack scenarios and identify attack paths, hotspots and potential mitigations. Consequently, predictive vulnerability management should not merely rank individual CVEs but should also consider their position within possible attack paths.

Empirical Review

Research by Jacobs et al. (2021) demonstrated that EPSS could provide probabilistic estimates of vulnerability exploitation and reported a ROC-AUC of approximately 0.838 in its original evaluation. This represented an important movement from severity-based assessment towards empirical exploitation prediction.

Peng et al. (2020) subsequently demonstrated that machine-learning-based remediation strategies could outperform conventional heuristic approaches because machine-learning models can incorporate interactions among multiple vulnerability characteristics.

Hemberg et al. (2022) developed a context-aware vulnerability prioritisation model that incorporated temporal characteristics and organisational context, addressing a major limitation of approaches that treat vulnerabilities independently of their environment. More recent research continues this trajectory. The 2026 Vulnerability Management Chaining framework integrated KEV, EPSS and CVSS into a sequential prioritisation model and reported substantial reductions in remediation workload while retaining high coverage of confirmed exploited vulnerabilities.

The literature nevertheless reveals three important gaps. First, many approaches concentrate on predicting vulnerability exploitation without adequately incorporating asset criticality. Second, vulnerability prioritisation is often separated from attack-path analysis. Third, relatively little empirical work integrates exploitation probability, infrastructure context and attack-path relevance into one predictive remediation framework. This study addresses these gaps through a controlled experimental AI framework.

3. Methodology

Research Design and Synthetic Dataset

This study adopted a simulation-based experimental research design to evaluate AI-driven vulnerability prioritisation. A synthetic dataset containing 20,000 vulnerability records was generated to represent CVE-level observations within a critical-

ISSN:3141-5571
E-ISSN:3141-558X

infrastructure environment. Each observation contained vulnerability and organisational-context variables designed to reproduce the principal factors relevant to exploitation prediction. The dependent variable was whether the vulnerability was classified as exploited within the target observation window.

The dataset contained CVSS severity, EPSS probability, KEV membership, public exploit availability, asset criticality, internet exposure, attack-surface level, asset age, patch complexity, dependency count, attacker interest and mitigation strength. These variables were selected because vulnerability exploitation depends not only on technical severity but also on exploit availability, threat attention and environmental context. The synthetic design allowed controlled experimentation without exposing confidential organisational vulnerability information.

The dataset was divided using stratified sampling into 70% training data, 15% validation data and 15% independent testing data. Random Forest, Support Vector Machine, XGBoost and LightGBM were trained using the training set and evaluated on the independent test set. Performance was assessed using accuracy, precision, recall, F1-score, ROC-AUC and false-positive rate. In addition, three conventional prioritisation baselines were evaluated: CVSS $\geq$ 7, EPSS $\geq$ 0.10 and KEV membership. A context-aware risk score was subsequently calculated by combining exploitation probability with confirmed exploitation, vulnerability severity, asset criticality and internet exposure.

Predictive Prioritisation Framework

The proposed framework consists of four stages. First, vulnerability intelligence is collected from CVE/CVSS, EPSS and KEV sources. Second, the vulnerability record is enriched with organisational information concerning asset criticality, internet exposure, attack surface and mitigation conditions. Third, the machine-learning model estimates exploitation likelihood. Finally, the predicted exploitation likelihood is combined with infrastructure context to assign remediation priority.

For attack-path analysis, vulnerabilities affecting critical or internet-facing assets are mapped to relevant adversarial techniques using the MITRE ATT&CK knowledge base. This enables security teams to distinguish isolated vulnerabilities from vulnerabilities that could facilitate initial access, privilege escalation, lateral movement or access to mission-critical systems. The framework therefore treats remediation as a risk-ranking problem rather than simply a severity-ranking exercise.

4. Results

Comparative Performance of AI Models

The independent test dataset contained 3,000 observations. The results were:

Model	Accuracy	Precision	Recall	F1-score	ROC-AUC	FPR
Random Forest	78.30%	80.03%	**94.35%**	86.60%	0.8084	68.18%
SVM	78.20%	79.87%	94.48%	86.57%	0.7708	68.96%
XGBoost	78.47%	**81.33%**	92.20%	86.42%	**0.8141**	**61.30%**
LightGBM	**78.50%**	81.19%	92.51%	**86.48%**	0.8090	62.08%

LightGBM achieved the highest accuracy at 78.50%, while XGBoost recorded the highest ROC-AUC at 0.8141 and precision at 81.33%. Random Forest achieved the highest recall at 94.35%. The relatively high recall across the models is particularly important for vulnerability management because failure to identify an exploitable vulnerability can have substantially greater consequences than generating an additional security alert.

ISSN:3141-5571
E-ISSN:3141-558X

The results also demonstrate why accuracy alone should not determine model selection. Although the models produced comparable accuracy, their false-positive rates remained relatively high. This indicates that vulnerability prioritisation is a more difficult task than binary vulnerability classification because many vulnerabilities possess characteristics associated with exploitation without necessarily becoming exploited.

Comparison with Conventional Prioritisation

Approach	Accuracy	Precision	Recall	F1-score
CVSS ≥ 7	60.73%	88.85%	53.95%	67.13%
EPSS ≥ 0.10	58.03%	76.93%	62.20%	68.78%
KEV membership	32.20%	91.88%	9.64%	17.45%

The comparison demonstrates the principal advantage of predictive modelling. CVSS-only classification achieved 60.73% accuracy and detected only 53.95% of the exploited observations. EPSS thresholding increased recall to 62.20%, while KEV membership produced very high precision but detected only 9.64% of exploited observations in the synthetic experiment.

The AI models produced considerably higher recall and F1-scores than the three conventional baselines. This finding supports the argument that vulnerability prioritisation should combine multiple signals rather than depend exclusively on severity, exploitation probability thresholding or confirmed exploitation status.

Context-Aware Remediation Prioritisation

The context-aware risk model classified the 3,000 test observations as follows:

Priority	Number	Percentage
Critical	1	0.03%
High	79	2.63%
Monitor	1,519	50.63%
Routine	1,401	46.70%
Total	**3,000**	**100%**

The classification demonstrates how contextualisation can reduce the immediate remediation workload. Rather than treating all vulnerabilities with similar technical severity as equally urgent, the framework identifies a smaller subset requiring immediate attention while assigning lower-risk vulnerabilities to monitoring or routine remediation cycles.

Attack-Path Implications

The proposed framework incorporates MITRE ATT&CK to determine whether high-risk vulnerabilities could contribute to plausible adversarial attack paths. Vulnerabilities associated with internet-facing assets receive additional attention because they may provide initial access. Where a vulnerable asset is connected to systems containing sensitive information or controlling critical operational processes, its priority increases even when its standalone CVSS score is not the highest.

This approach is consistent with MITRE's critical-infrastructure modelling approach, which demonstrates how CVEs can be linked to ATT&CK techniques and infrastructure attack paths.

ISSN:3141-5571
E-ISSN:3141-558X

5. Discussion

The empirical results support the central proposition of the study: AI can improve vulnerability prioritisation by learning from multiple vulnerability and organisational characteristics simultaneously. The superiority of the machine-learning models over CVSS-only, EPSS-threshold and KEV-only baselines is consistent with earlier research showing that machine learning can incorporate interactions that simple heuristic approaches cannot adequately represent (Peng et al., 2020).

However, the results also reveal an important practical limitation. The relatively high false-positive rates demonstrate that exploitation prediction remains probabilistic. AI should therefore not be interpreted as replacing cybersecurity analysts. Instead, it should function as a decision-support mechanism that reduces the remediation search space and directs expert attention toward vulnerabilities with the strongest combination of exploitation likelihood and organisational consequence.

The findings also reinforce the importance of contextual vulnerability management. EPSS answers the question of whether a vulnerability is likely to be exploited, while asset criticality addresses what happens if exploitation succeeds. Attack-path analysis adds a third dimension by asking what an attacker could reach after exploiting the vulnerability. Combining these dimensions provides a substantially richer basis for remediation decisions.

The framework therefore represents a movement from conventional vulnerability management towards predictive, contextual and attack-path-aware vulnerability management. This is particularly relevant to critical infrastructure, where patching every vulnerability immediately is operationally unrealistic and, in some OT environments, potentially unsafe.

6. Conclusion and Recommendations

Conclusion

This study developed and empirically evaluated an AI-based predictive vulnerability management framework for critical infrastructure cybersecurity. The study demonstrated that machine-learning models can integrate vulnerability characteristics, exploitation indicators and organisational context to improve the identification of vulnerabilities likely to be exploited.

Among the evaluated algorithms, LightGBM achieved the highest accuracy at 78.50%, while XGBoost achieved the highest ROC-AUC of 0.8141. The AI models also substantially outperformed conventional CVSS, EPSS-threshold and KEV-only baselines in balanced predictive performance. The findings demonstrate that vulnerability severity alone provides an incomplete basis for remediation prioritisation.

The study further established that effective vulnerability management should extend beyond individual CVE ranking. Incorporating asset criticality, internet exposure and attack-path relevance enables organisations to identify vulnerabilities that could provide strategically important routes into critical infrastructure. The proposed framework therefore provides a foundation for moving from reactive vulnerability management towards predictive and risk-informed cybersecurity decision-making.

Recommendations

Organisations should integrate AI into vulnerability management programmes to rank vulnerabilities according to predicted exploitation probability rather than relying exclusively on CVSS severity.

EPSS, CVSS and KEV should be used together rather than independently. EPSS provides exploitation probability, CVSS provides technical severity, while KEV provides evidence of observed exploitation. Asset criticality should be incorporated into remediation decisions. Vulnerabilities affecting mission-critical or safety-sensitive systems should receive greater attention than equivalent vulnerabilities affecting low-value assets.

Internet-facing assets should receive additional priority, particularly where vulnerabilities can provide initial access to critical infrastructure. MITRE ATT&CK should be integrated with vulnerability-management platforms to identify vulnerabilities that could form part of realistic attack paths.

Organisations should adopt risk-based remediation queues rather than attempting to patch vulnerabilities strictly according to their CVSS ranking.

Human cybersecurity expertise should remain part of the decision process. AI predictions should support, rather than replace, security analysts, particularly where patching could affect operational availability or safety.

ISSN:3141-5571
E-ISSN:3141-558X

Limitations

The principal limitation is that the empirical dataset was synthetically generated rather than obtained from a live critical-infrastructure organisation. The findings should therefore be interpreted as evidence from a controlled computational experiment rather than direct evidence of operational performance. The simulation also simplifies complex relationships among vulnerabilities, assets, attackers and network architecture. Consequently, future research should validate the framework using real enterprise vulnerability data, historical exploitation records and actual infrastructure attack graphs.

7. References

Bulut, M. F., Adebayo, A., Sow, D., & Ocepek, S. (2022). *Vulnerability prioritization: An offensive security approach*. arXiv. https://arxiv.org/abs/2206.11182

Cybersecurity and Infrastructure Security Agency. (2021). *Reducing the significant risk of known exploited vulnerabilities*. U.S. Department of Homeland Security.

Cybersecurity and Infrastructure Security Agency. (2022). *Federal government cybersecurity incident and vulnerability response playbooks*. U.S. Department of Homeland Security.

Cybersecurity and Infrastructure Security Agency. (2024). *Securing the software supply chain: Recommended practices for software bill of materials consumption*. U.S. Department of Homeland Security.

Hemberg, E., Zipkin, J., Skowyra, R., & others. (2022). CAVP: A context-aware vulnerability prioritization model. *Computers & Security, 116*, 102639. https://doi.org/10.1016/j.cose.2022.102639

Jacobs, J., Romanosky, S., Edwards, B., Adjerid, I., & Roytman, M. (2021). Exploit Prediction Scoring System (EPSS). *Digital Threats: Research and Practice, 2*(3), Article 20. https://doi.org/10.1145/3436242

Peng, J., Malin, B., & others. (2020). Improving vulnerability remediation through better exploit prediction. *Journal of Cybersecurity, 6*(1), tyaa015. https://doi.org/10.1093/cybsec/tyaa015

MITRE. (2022). *Critical Infrastructure Cyberspace Analysis Tool (CICAT)*. MITRE Corporation.

MITRE. (2026). *MITRE ATT&CK: Enterprise matrix*. MITRE Corporation.

FIRST. (2026). *Exploit Prediction Scoring System (EPSS)*. Forum of Incident Response and Security Teams.

Vulnerability Management Chaining: An integrated framework for efficient cybersecurity risk prioritization. (2026). *IEEE Access, 14*, 31407–31424. https://doi.org/10.1109/ACCESS.2026.3665768