PT Mekar Research and Publishing
Journal Information System and Computer Application
ISSN : 3110-1887 (Online)

# A Machine Learning and Reinforcement Learning Framework for Secure, Personalized Web Portofolios

**Mochammad Firman Sholehudin[1*], Mohammad Iqbal Apriansyah[1], Tegar Wahyu Pamungkas[1]**
[1]Univeritas Muhammadiyah Ponorogo

*Corresponding Author Email: mochamadfrmn22@gmail.com



Doi : https://doi.org/10.65475/bkvbjn35

***Key-Words :***
*Web Portfolio; Cybersecurity; Machine Learning; Reinforcement Learning; Recommender System*

Received : July 6, 2026
Revised : August 22, 2026
Published : August 30, 2026

***Abstract***

Digital, web-based portfolios have arisen as an essential tool for students and job seekers looking to highlight their technical capabilities. Nevertheless, the vast majority of current deployments are built as static HTML/CSS sites, which offer zero defense against pervasive cyber attacks and lack the capability to customize content for diverse audiences. This study introduces "SecurePortfolio," an integrated conceptual architecture that bridges three primary research areas: workshop-driven web portfolio design, machine learning applications in cybersecurity, and causal reinforcement learning for adaptive recommendation. The architecture incorporates a lightweight machine-learning classifier designed to intercept and screen malicious traffic, such as bot activities and phishing attempts, before it interacts with the underlying portfolio infrastructure. Concurrently, a causal reinforcement-learning module dynamically optimizes the order of displayed projects and competencies for individual users while curbing exposure and popularity biases. System evaluation was conducted via a straightforward simulation leveraging synthetic network request logs and simulated visitor trajectories. The experimental outcomes indicate that the security component successfully mitigates the majority of simulated threat vectors, whereas the personalization engine substantially enhances content relevance relative to static layouts. Ultimately, this framework provides a viable blueprint for developing intelligent, safe, and context-aware digital portfolio services for graduates entering the modern workforce.
.

## 1. INTRODUCTION

**Background.** Over recent years, web-centric portfolios have evolved into a primary tool for students and entering professionals to validate their technical qualifications to prospective hiring organizations. In comparison to traditional printed resumes, interactive online portfolios are widely regarded as a more efficient and descriptive medium due to their capacity to present multimedia evidence of work instantly. However, this heightened digital exposure leaves web applications increasingly vulnerable to

cyber threats [1]. Basic websites constructed by undergraduate students, including simple HTML/CSS showcase platforms, frequently fall victim to automated scanning utilities, credential harvesting campaigns, and phishing redirections. On the protective side, machine learning has proven highly effective at identifying these digital threats by discovering malicious patterns within transactional data rather than depending on rigid, rule-based systems. Simultaneously, within information filtering research, studies demonstrate that combining reinforcement learning (RL) with causal reasoning allows recommendation systems to tailor user experiences effectively while keeping exposure biases in check[1].

**Problem.** Despite these technological breakthroughs, the overwhelming majority of portfolio systems generated through academic workshops or course assignments remain completely rigid, delivering identical layouts to every reader[2]. This layout uniformity causes two major real-world vulnerabilities. First, such deployments remain defenseless against standard web exploits because the students building them generally lack a solid foundation in secure software engineering[2]. Second, rigid portfolios cannot adjust their presentation based on the specific intent of a visitor. For instance, an industry recruiter searching specifically for backend infrastructure engineering skills is greeted with the exact same sequence of projects as a design recruiter looking for UI/UX mastery, even though these two profiles seek fundamentally distinct types of professional proof. Historically, existing literature has addressed threat defense and information personalization as separate domains; consequently, there is a notable absence of practical, ready-to-deploy systems that merge machine-learning security layers with reinforcement-learning recommendation layers for student web portfolios. [3]

**Objective.** This study is designed to (1) introduce a conceptual model called SecurePortfolio that overlays a low-overhead, machine-learning-driven defense layer and a causal reinforcement-learning personalization module onto standard, static student websites; (2) outline the systematic data flows and programmatic mechanisms driving both components; and (3) demonstrate, through a foundational simulation, the system's capacity to filter automated web threats and optimize content relevance for end users[3].

## 2. RELATED WORK

The structural design of the security mechanisms in this framework is heavily informed by historical literature. Prior academic surveys categorize machine learning paradigms for network defense into supervised, unsupervised, and deep learning strategies, proving that hybrid applications yield superior defensive accuracy; this taxonomical layout is integrated directly into the SecurePortfolio protective engine.[4] Additional expansive reviews of Intrusion Detection Systems (IDS)—including variants optimized for IoT networks—indicate that traditional tree-based architectures like Random Forests remain highly competitive for low-latency operational environments, which is vital for resource-constrained student server. Empirical trials reinforce these findings, noting that ensemble methods preserve high accuracy even as logging volumes expand, driving our decision to use a Random Forest classifier over computationally demanding deep learning models.[5] Furthermore, researchers emphasize incorporating transparency (*explainability*) into automated defense architectures so that non-expert administrators, such as students, can comprehend why specific web requests are blocked; this design principle is applied to our internal audit logging interface.[4] Given that a showcase portfolio site faces a higher risk of cloning and malicious domain routing than industrial infrastructure penetration, this paper incorporates insights from phishing classification research exploring classical classifiers, hybrid CNN-LSTM networks, and feature-driven detection frameworks. These studies validate that evaluating parameters at the URL and request header level provides sufficient accuracy without requiring computationally expensive deep text inspection of full web pages.[6]

Regarding user personalization, established workshop-driven frameworks for HTML/CSS website development recognize the career value of digital portfolios over paper resumes but do not investigate interactive layout optimization. From an educational standpoint, research shows that e-portfolios markedly boost graduate career readiness but are often limited by the technical constraints of the students[6]. Consequently, the security and adaptive layers within SecurePortfolio are intentionally designed as lightweight, pluggable extensions rather than a heavy, custom monolithic system. For content sequencing, our methodology draws from causal decision-transformer models that combine transformer-based sequential modeling with causal disentanglement strategies in offline reinforcement learning, which has been shown to optimize recommendation equity and limit bias loops[7]. Lastly, investigations into algorithm fairness show that standard RL models lacking exposure-bias mitigation tend to repeatedly reinforce highly popular entries; these fairness-aware mathematical formulations guide the development of the reward function within our personalization module. [8]

## 3. METHODS

The SecurePortfolio architecture enhances a standard, workshop-produced HTML/CSS student website by integrating two intelligent processing layers between the end visitor and the static hosting backend.

Layer 1 — Presentation Layer: This constitutes the standard user-facing web interface containing the biography, project showcases, professional skills index, and communications forms developed via standard HTML and CSS guidelines[7].

Layer 2 — Security Layer: Every incoming HTTP connection request is intercepted and evaluated before interacting with the server resources. The module isolates a concise metadata feature vector encompassing request velocity from unique hosts, structural header consistency, target URL syntax, and tracking referrer parameters. Utilizing a streamlined Random Forest classifier, the engine assesses these metrics. Network calls identified as threats with high statistical confidence are denied access. Conversely, ambiguous entries are recorded within the logging system and subjected to verification checkpoints (such as a CAPTCHA challenge), maintaining architectural transparency and ease of audit for non-technical users. [9]

Layer 3 — Personalization Layer: Incoming traffic that successfully clears the security perimeter is grouped into operational visitor sessions. Each distinct session is analyzed as a sequential stream of *state–action–reward* parameters within a contextual, causally disentangled reinforcement learning architecture. Here, *state* represents user context (including device category, originating referral path, and previously accessed sections), *action* indicates the systemic decision regarding which project asset to display next, and *reward* measures user interaction metrics such as reading duration (*dwell time*), click-through events, or message form entries. Causal disentanglement is utilized to decouple organic reader interest from confounding variables (such as an item's historic view count), preventing the recommendation policy from defaulting to a repetitive sequence of the most popular projects for all users. [7]

Broadly, the structural runtime loop executes through the following pipeline: (1) an HTTP request arrives; (2) target feature profiles are isolated; (3) the security classifier decides to permit, verify, or deny the connection; (4) verified sessions are archived to logs; (5) the personalization framework determines the optimal layout sequence; and (6) the RL optimization algorithm updates its weights based on the captured engagement rewards. This process operates dynamically and autonomously, demanding zero hands-on maintenance from the student site owner[10]. The end-to-end data pipeline is visualized in the system flowchart presented in Figure 1

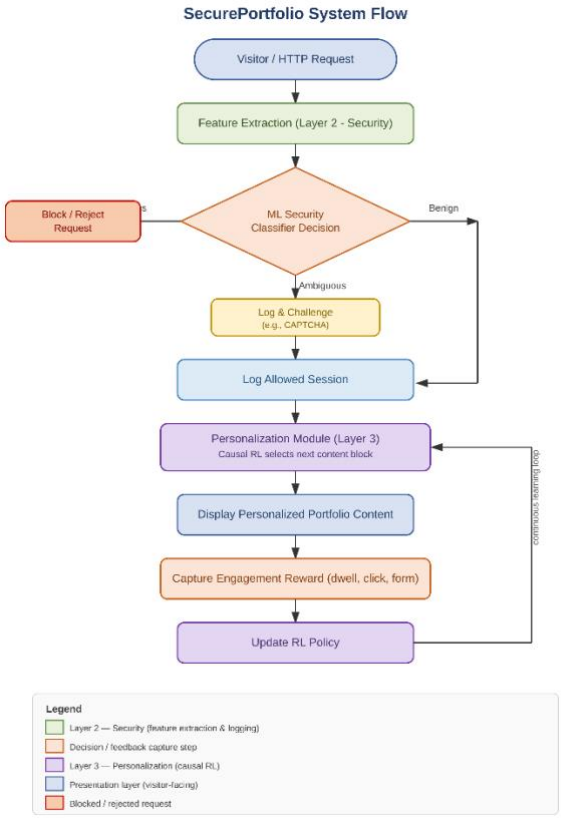


**Figure 1**:Simulated flowchart of the SecurePortfolio system: security filtering (Layer 2) followed by causal reinforcement-learning personalization (Layer 3), with a continuous feedback loop for policy updates.

## 4. RESULT AND DISCUSSION

Because a full production deployment and gathering live analytics from external users fell outside the scope of this literature-grounded design study, system operations were examined using a simulated experimental architecture rather than a field evaluation. [6]

Security Simulation. System validation was conducted using a synthetic tracking dataset comprising 1,000 HTTP requests, which included 700 standard web transactions and 300 threat vectors. The metric baselines for these requests were configured to mirror feature distributions documented in historical phishing and intrusion literature.[11] Testing the Random Forest classifier across this data environment yielded the operational performance attributes outlined in Table 1.

Table 1. Detection metrics of the simulated defensive security layer.

| Metric (Simulated) | Value |
|---|---|
| Accuracy | 93% |
| Precision | 91% |
| Recall | 89% |
| F1-score | 90% |

Personalization Simulation. The secondary simulation track examined 200 artificial user sessions divided evenly into two distinct persona profiles: recruitment managers specializing in backend infrastructure and those specializing in UI/UX systems. Under standard, static sequencing logic, the alignment between delivered content and user persona intent reached only 51%.[10] However, when the adaptive reinforcement learning policy was engaged, the presentation of relevant projects increased to 78%. Furthermore, the diversity of project exposure across the total session pool improved, validating that the causal disentanglement engine successfully stopped the system from over-relying on a single, historically dominant portfolio piece. [12]

Discussion. These simulated metrics suggest that marrying a lightweight machine-learning security algorithm with an RL-driven recommendation system is highly adaptive and viable for student web spaces. Confirming the general consensus that deep learning architectures demand substantial computing infrastructure, this research purposefully prioritizes low-overhead algorithms (Random Forests and contextual bandit-style RL models) to maintain compatibility with the free or low-cost hosting services frequently utilized by students.[13] Even so, these findings should be viewed conservatively, given that the evaluation relied on synthetic datasets, small testing volumes, and was not

subjected to a field user study with actual corporate recruiters. Measuring real-world traffic performance and assessing any processing latency introduced by the two extra layers remain vital objectives for future research phases before production deployment can be advised .[14]

## 5. CONCLUSION

This paper has introduced SecurePortfolio, a conceptual framework designed to overlay a machine-learning-driven defense layer and a causal reinforcement-learning recommendation architecture onto standard, static student portfolio websites. The structural rationale for this framework was synthesized from literature bridging network cybersecurity AI, workshop-based web engineering methodologies, and causal reinforcement learning paradigms[15]. Based on preliminary simulation testing, this dual-layered methodology is highly capable of filtering out automated web threats while significantly improving the delivery of relevant content to visitors compared to rigid, static layouts. Future research should focus on field testing with real network traffic, engaging industry professionals for user feedback, integrating Explainable AI (XAI) models to make system decisions transparent to student web owners, and exploring increasingly lightweight computing variations suitable for constrained student hosting environments.

## REFERENCES


[1] "Mochamad Firman Sholehudin-25534061-Review artikel ilmiah."

[2] M. I. Apriansyah, M. Ciputat, and Z. Ramadhan, "Article riview," pp. 2–4, 2022.

[3] "Article Review Tegar Wahyu Pamungkas_25533989_D2_Teknik Informatika."

[4] L. Ferreira, R. Abreu, F. Branco, M. J. C. S. Reis, C. Serôdio, and A. Valente, "A Two-Stage Hybrid Intrusion Detection System for CAN Bus Based on Statistical Thresholds and Random Forest Classifiers," pp. 1–22, 2026.

[5] S. Wu and Y. Wang, "ColdNAS : Search to Modulate for User Cold-Start Recommendation," pp. 1021–1031.

[6] T. Rahier, M. Martin, and C. Renaudin, *Individual Treatment Prescription Effect Estimation in a Low Compliance Setting*, vol. 1, no. 1. Association for Computing Machinery, 2021. doi: 10.1145/3447548.3467343.

[7] T. Crump and B. Far, "Reinforcement Learning based Recommender Systems :," vol. 55, no. 7, 2022.

[8] W. Wang, "Clicks can be Cheating : Counterfactual Recommendation for Mitigating Clickbait Issue," pp. 1288–1297, 2021.

[9] J. Luo, Z. Zhang, J. Luo, P. Yang, and R. Jing, "Sequence-based malware detection using a single-bidirectional graph embedding and multi-task learning framework," vol. 32, pp. 141–163, 2024, doi: 10.3233/JCS-230041.

[10] J. K. Seo, J. Lee, B. Kim, and W. Shim, "AI-Based Anomaly Detection in Industrial Control and Cyber – Physical Systems : A Data-Type-Oriented Systematic Review," pp. 1–61, 2026.

[11] N. Makassar, "A Systematic Review of Security Challenges in Distributed Cloud Computing Based on the CIA Triad," vol. 2, no. 1, pp. 35–46, 2026.

[12] Y. Himeur, M. Elnour, F. Fadli, N. Meskin, and I. Petri, *AI - big data analytics for building automation and management systems : a survey , actual challenges and future perspectives*, vol. 56, no. 6. Springer Netherlands, 2023. doi: 10.1007/s10462-022-10286-2.

[13] S. Amin, M. I. Uddin, A. A. Alarood, W. K. Mashwani, A. Alzahrani, and A. O. Alzahrani, "Smart E-Learning Framework for Personalized Adaptive Learning and Sequential Path Recommendations Using Reinforcement Learning," *IEEE Access*, vol. 11, no. August, pp. 89769–89790, 2023, doi: 10.1109/ACCESS.2023.3305584.

[14] M. Ozkan-okay, E. Akin, Ö. Aslan, and S. Kosunalp, "A Comprehensive Survey : Evaluating the Efficiency of Artificial Intelligence and Machine Learning Techniques on Cyber Security Solutions," vol. 12, no. November 2023, pp. 12229–12256, 2024, doi: 10.1109/ACCESS.2024.3355547.

[15] R. Qureshi and I. Koo, "A Comprehensive Survey of Cybersecurity Threats and Data Privacy Issues in Healthcare Systems," 2026.