

Enhancing CI/CD Security with Artificial Intelligence: State of the Art, Challenges, and Integrated Approaches

Iveri Jajanidze

PhD Candidate (Informatics), Doctoral Program in Informatics, Georgian Technical University, Tbilisi, Georgia.

Email: jajanidze.iveri24@gtu.ge

Abstract

Continuous Integration and Continuous Delivery (CI/CD) pipelines have become a core component of modern software engineering, enabling rapid and automated deployment of applications. While these practices significantly improve development efficiency, they also introduce new and complex security risks, including supply chain attacks, configuration drift, and large-scale propagation of vulnerabilities through automated pipelines [6, 7, 8]. Traditional security mechanisms, which rely primarily on static rules and signature-based detection, are increasingly inadequate in highly dynamic DevSecOps environments [3, 12]. This paper investigates the application of Artificial Intelligence (AI) and Machine Learning (ML) techniques to enhance the security of CI/CD systems. The proposed approach integrates anomaly detection, supervised classification, Explainable AI (XAI), federated learning, and adversarial robustness mechanisms into a unified analytical architecture [2, 4, 5]. The study argues that explainability and privacy-preserving learning are critical for adoption in regulated and mission-critical environments [2, 4, 7]. A conceptual and architectural framework is presented, together with an experimental evaluation strategy, demonstrating how AI-driven security analytics can provide proactive, transparent, and scalable protection for modern CI/CD pipelines.

Keywords: CI/CD Security, DevSecOps, Artificial Intelligence, Machine Learning, Explainable Artificial Intelligence, Federated Learning, Adversarial Machine Learning

1. Introduction

Modern software development is increasingly dependent on Continuous Integration (CI) and Continuous Delivery/Deployment (CD) practices, which enable frequent code integration, automated testing, and rapid release cycles [12, 13]. These approaches form the foundation of DevSecOps, where development, operations, and security are expected to function as a unified and highly automated process [12]. While CI/CD pipelines significantly improve delivery speed and operational efficiency, they also expand the attack surface by introducing complex dependencies, extensive automation, and a high frequency of changes across infrastructure and application layers [6, 7].

A defining characteristic of CI/CD environments is the large-scale use of Infrastructure as Code (IaC), third-party libraries, containerization platforms, and automated orchestration tools [12, 13]. In such settings, a single misconfiguration, compromised dependency, or malicious code injection can propagate through the entire pipeline in a very short time, potentially resulting in large-scale security incidents, including supply chain attacks [3, 9]. Recent incidents and regulatory responses clearly indicate that pipeline security is no longer a peripheral concern but a central requirement for trustworthy software delivery [6, 7, 8].

Conventional security controls in software pipelines are typically based on static analysis, predefined rules, and signature-based detection mechanisms. Although these techniques remain useful, they struggle to cope with the dynamic and heterogeneous nature of modern CI/CD systems and often fail to detect previously unseen attack patterns [3, 12]. As a result, security must be embedded into the pipeline as an adaptive, data-driven, and continuously improving process rather than as a final checkpoint before deployment [12, 13].

In this context, Artificial Intelligence (AI) and Machine Learning (ML) provide promising capabilities for analyzing large volumes of heterogeneous data generated by CI/CD systems, including source code changes, build and deployment logs, configuration files, and runtime telemetry [3, 12]. AI-driven methods enable automated anomaly detection, risk classification,

and behavioral analysis at a scale and speed that is not feasible with purely manual or rule-based approaches [3]. At the same time, modern threats increasingly leverage AI themselves, such as adaptive malware and automatically generated social engineering content, which further strengthens the need for AI-assisted defensive mechanisms [3, 11].

However, the adoption of AI in security-critical and regulated environments raises additional requirements. Decision transparency and auditability are essential, particularly in sectors subject to strict compliance frameworks such as ISO/IEC 27001:2022, NIST SP 800-53, and the European Union’s Digital Operational Resilience Act (DORA) [6, 7, 8]. Explainable Artificial Intelligence (XAI) addresses this challenge by providing human-interpretable explanations of model outputs, making automated decisions more trustworthy and operationally useful [2, 5].

At the same time, data protection and regulatory constraints often limit the feasibility of centralizing sensitive operational data for model training. Federated learning offers a privacy-preserving alternative by enabling collaborative model training across multiple environments without sharing raw data [4]. Another important challenge is the growing relevance of adversarial attacks against ML models, including data poisoning and evasion techniques, which can be particularly dangerous in highly automated CI/CD pipelines [5, 12].

This paper presents an integrated approach to CI/CD security that combines anomaly detection, supervised classification, XAI, federated learning, and adversarial modeling into a unified architectural framework [2, 3, 4, 5]. The objective is to demonstrate how AI-driven security analytics can move CI/CD protection from a predominantly reactive posture to a proactive, explainable, and privacy-aware security model suitable for modern DevSecOps ecosystems [6, 7, 8].

2. Related Work and Problem Formulation

The security of CI/CD and DevSecOps pipelines has attracted increasing attention in recent years due to the growing number of supply chain attacks, configuration-based vulnerabilities, and large-scale automation failures [3, 9, 12]. Existing research and industrial practices emphasize the importance of integrating security controls directly into development and deployment workflows, rather than treating security as a separate, post-deployment activity [12, 13]. Standards and regulatory frameworks such as NIST SP 800-53, ISO/IEC 27001:2022, and the European Union’s DORA further reinforce the need for systematic, auditable, and resilient security mechanisms in software delivery infrastructures [6, 7, 8].

A substantial body of work has explored the application of machine learning techniques for security monitoring, including malware detection, intrusion detection, and anomaly detection in operational data [3]. Surveys of ML-based malware analysis show that supervised and unsupervised learning methods can outperform traditional signature-based approaches, particularly in detecting previously unseen threats [3]. In CI/CD contexts, these techniques are increasingly applied to analyze build logs, configuration changes, and runtime telemetry in order to identify abnormal behavior patterns that may indicate compromise or misconfiguration [12, 13].

Explainable Artificial Intelligence has emerged as a critical research direction for safety- and security-critical applications [2, 5]. Methods such as LIME and SHAP have been proposed to provide local and global explanations of model decisions, enabling better human understanding and auditability of automated systems [2, 5, 17, 18]. In security operations, such transparency is essential for incident investigation, compliance verification, and trust in automated decision-making processes [2, 5]. Nevertheless, many existing ML-based security tools still operate as “black boxes,” which limits their acceptance in regulated environments [2, 5].

Another important research direction addresses data privacy and governance in distributed infrastructures. Federated learning has been proposed as a solution that enables collaborative model training across multiple environments without sharing raw data [4]. Prior

studies demonstrate that federated learning can preserve model performance while significantly reducing the risk of sensitive data exposure, which is particularly relevant for CI/CD ecosystems spanning multiple organizational or regulatory domains [4, 7, 8].

At the same time, the robustness of machine learning models against adversarial manipulation has become a major concern. Adversarial examples, data poisoning, and evasion attacks show that ML models can be intentionally misled with carefully crafted inputs or corrupted training data [5]. In highly automated CI/CD pipelines, such weaknesses are especially dangerous, because erroneous model decisions can propagate rapidly and affect large parts of the delivery process [12]. Although adversarial training and robustness optimization techniques can mitigate some risks, they do not provide a complete solution and require continuous evaluation and adaptation [5, 6].

Despite these advances, several gaps remain. Many existing solutions address only isolated aspects of CI/CD security without integrating explainability, privacy preservation, and robustness into a single coherent framework [2, 3, 4, 5]. Moreover, the operational integration of AI-based security analytics into CI/CD pipelines is often insufficiently addressed, particularly with respect to automated decision points and human-in-the-loop oversight [12, 13]. Finally, regulatory and compliance requirements are rarely treated as first-class design constraints in technical security architectures [6, 7, 8].

Based on these observations, the problem addressed in this paper can be formulated as follows: how to design an AI-driven security architecture for CI/CD pipelines that is effective in detecting complex threats, transparent in its decision-making, respectful of data privacy constraints, and resilient against adversarial manipulation. The proposed approach aims to fill this gap by combining anomaly detection, supervised classification, Explainable AI, federated learning, and adversarial modeling into an integrated DevSecOps-oriented framework, which is detailed in the following sections.

3. Methodology and Core Models

The proposed approach is based on the integration of several complementary analytical components that address different aspects of CI/CD security. The methodology combines data-driven anomaly detection, supervised classification, explainable decision support, privacy-preserving collaborative learning, and adversarial robustness mechanisms within a unified framework [2, 3, 4, 5]. This design follows the general principles of secure and auditable systems emphasized in contemporary standards and regulatory frameworks [6, 7, 8].

A. Data Sources and Preprocessing

Security analysis in CI/CD environments relies on heterogeneous data originating from multiple sources, including version control systems, pipeline execution logs, Infrastructure as Code (IaC) configurations, container metadata, and authentication or network events [12, 13]. These data streams are characterized by high volume, partial structure, and strong temporal dynamics, which necessitates a systematic preprocessing stage [12]. This stage includes normalization, deduplication, temporal aggregation, and feature extraction, as commonly recommended in ML-based security analytics [3, 12].

Textual data, such as commit messages and log entries, are processed using natural language processing techniques, including tokenization and vectorization based on term-frequency or embedding representations [12]. The resulting feature vectors form the input to the machine learning models employed in subsequent stages, enabling both statistical learning and behavioral analysis across the CI/CD pipeline [12].

B. Supervised and Unsupervised Learning

Two complementary learning paradigms are used to address different threat detection scenarios. Supervised learning models are applied when labeled examples of known threats are available, following established approaches in malware analysis and security event classification

[3]. Typical classifiers include logistic regression, decision trees, gradient boosting methods such as XGBoost, and neural networks, which have demonstrated strong performance in security-related classification tasks [3, 12].

In contrast, unsupervised learning is primarily employed for the detection of previously unseen or rare events, which is particularly important in dynamic CI/CD environments where new failure modes and attack patterns continuously emerge [3, 12]. Autoencoders and isolation-based methods are widely used for this purpose. Autoencoders are trained to reconstruct normal behavior patterns, and deviations from this learned representation are interpreted as potential anomalies [12]. A high reconstruction error therefore serves as an indicator of abnormal or suspicious activity, which has proven effective in large-scale security monitoring scenarios [12].

C. Explainable Artificial Intelligence

In security-critical contexts, the interpretability of automated decisions is as important as their predictive accuracy [2, 5]. For this reason, the framework integrates Explainable Artificial Intelligence (XAI) techniques that provide human-understandable explanations for model outputs. Model-agnostic methods such as LIME and SHAP are used to estimate the contribution of individual features to a specific prediction and to explain both local and global model behavior [2, 5, 17, 18].

These explanations enable security analysts to verify whether a decision is based on meaningful indicators, to support incident investigation, and to satisfy audit and compliance requirements [2, 5]. By embedding explainability directly into the analytical pipeline, the system addresses a major limitation of many ML-based security tools, which often operate as opaque “black boxes” and therefore face resistance in regulated or mission-critical environments [2, 5].

D. Federated Learning and Privacy Preservation

In many organizational settings, CI/CD data cannot be centralized due to regulatory, contractual, or operational constraints [4, 7, 8]. To address this limitation, the proposed

methodology incorporates federated learning as a core training paradigm. In this setup, models are trained locally within separate environments, such as development, testing, and production domains, and only model updates are shared with a central coordinator [4]. The global model is obtained through aggregation of these local updates, which reduces the risk of sensitive data exposure while preserving the benefits of collaborative learning [4].

This approach is particularly suitable for CI/CD ecosystems spanning multiple administrative or regulatory domains, where data governance requirements make traditional centralized training impractical [4, 7, 8].

E. Adversarial Robustness Considerations

The increasing use of machine learning in security systems has introduced new attack vectors, commonly referred to as adversarial attacks [5]. These include adversarial examples, data poisoning, and evasion techniques that aim to mislead models during inference or corrupt them during training [5]. In highly automated CI/CD pipelines, such weaknesses are especially critical, as erroneous model decisions can propagate rapidly through deployment workflows [12].

To mitigate these risks, the proposed framework considers adversarial training and robustness-oriented evaluation as integral parts of the model lifecycle [5]. Although no defense can guarantee complete protection, systematic robustness testing and continuous model adaptation significantly improve resilience against realistic attack scenarios [5, 6].

4. System Architecture and CI/CD Integration

The proposed security analytics system is designed as a modular, multi-layer architecture that supports scalable data collection, automated analysis, explainable decision-making, and operational integration within CI/CD environments [12, 16]. The architecture follows the principles of separation of concerns and microservice-based design in order to ensure

extensibility, maintainability, and adaptability to different organizational and technological contexts [12, 14, 15].

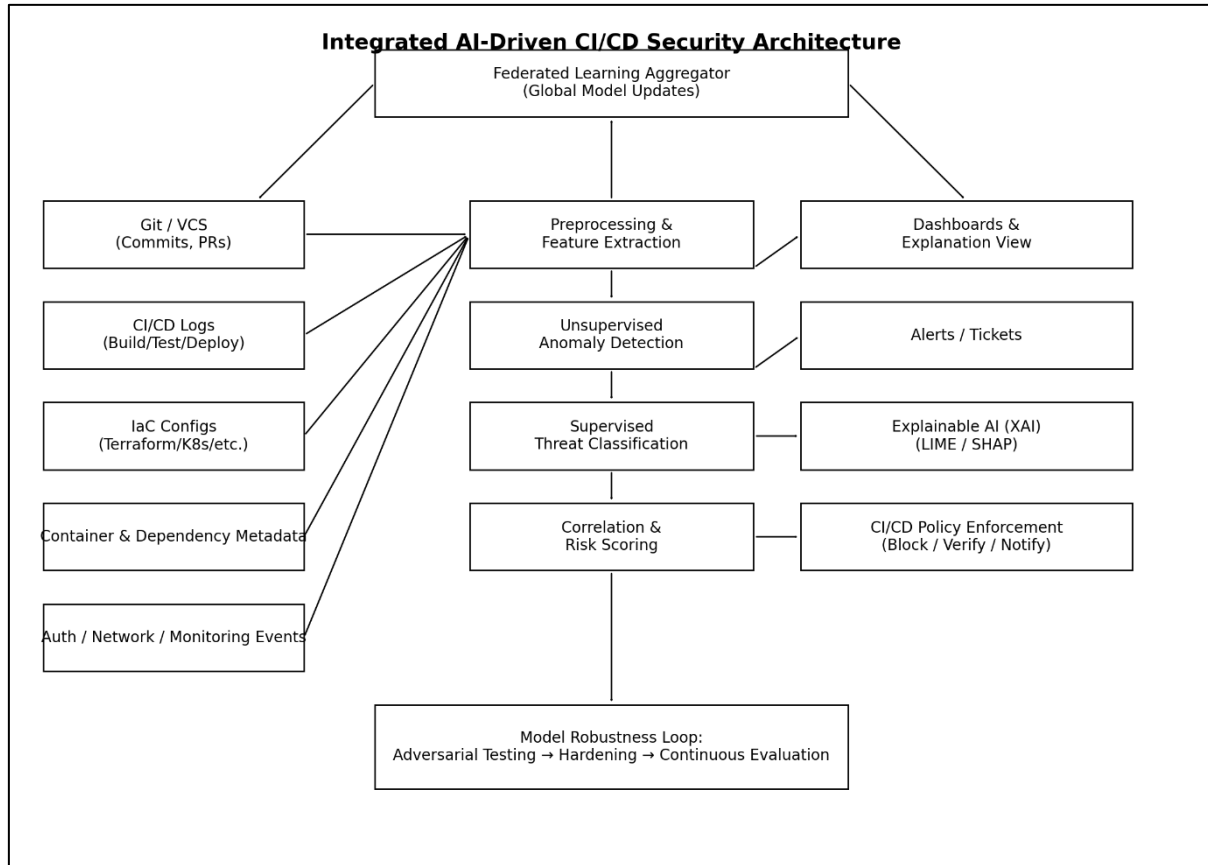
A. Overall Architecture

At the lowest layer, the system is responsible for collecting heterogeneous data from multiple CI/CD-related sources, including version control systems, CI/CD orchestrators, container platforms, and monitoring components [12, 13, 15, 16]. These data include source code changes, pipeline execution logs, Infrastructure as Code (IaC) configurations, container metadata, and security-relevant events [12, 13]. The collected information is normalized and forwarded to a central processing module, where preliminary filtering and feature preparation are performed [12].

The intermediate layer hosts the analytical components, including machine learning models for anomaly detection and classification, as well as Explainable AI modules for interpreting model outputs [2, 5, 12]. This layer constitutes the core of the decision-making process and is responsible for correlating events across different data sources in order to identify potential security incidents [3, 12]. In parallel, a federated learning service manages collaborative model updates across multiple environments, enabling continuous improvement of detection capabilities without violating data governance constraints [4, 7, 8].

The upper layer is dedicated to visualization and operational integration. It provides dashboards, metrics, and explanation views that support both real-time monitoring and forensic analysis of incidents [16]. In addition, this layer exposes interfaces for automated or semi-automated actions within the CI/CD pipeline, such as blocking a deployment, triggering additional verification steps, or notifying security personnel [12, 13]. Please refer to figure 1, Integrated AI-Driven CI/CD Security Architecture.

Figure 1 - Integrated AI-Driven CI/CD Security Architecture



B. Data Management and Storage

The system employs a centralized storage layer that supports both structured and semi-structured data, enabling efficient handling of logs, configuration files, and extracted feature vectors [12]. This storage layer maintains raw event data as well as processed representations used for model training and evaluation, which is essential for reproducibility and auditability of experimental results [6, 7].

Versioning of datasets and models is treated as a first-class requirement, following established MLOps practices for controlled experimentation and traceability [12]. This approach allows historical analyses to be repeated and supports compliance with regulatory expectations regarding accountability and evidence-based decision-making [6, 7, 8].

C. Integration into CI/CD Pipelines

A central design goal of the proposed architecture is seamless integration into existing CI/CD pipelines, so that security analysis becomes a natural and continuous part of the development and delivery process rather than an external or ad hoc activity [12, 13]. To achieve this, the system provides integration points at multiple stages of the pipeline, including code commit, build, test, and deployment phases [12, 14, 15].

At each stage, relevant data are collected and analyzed, and the results are fed back into the pipeline control logic. In critical cases, such as the detection of high-risk anomalies or policy violations, the pipeline can be automatically halted or redirected to additional validation steps [6, 7, 12]. In less severe scenarios, the system operates in an advisory mode, providing risk assessments and explanations to developers and operators while leaving the final decision to human oversight [12, 13].

This hybrid approach ensures a balance between automation and operational flexibility, which is essential for maintaining development velocity while enforcing robust security controls [12, 13].

D. Monitoring, Visualization, and Operational Use

Operational effectiveness requires that analytical results be presented in a clear and actionable form. For this purpose, the architecture integrates monitoring and visualization components that aggregate metrics, analyze time series, and present security-relevant information through dashboards and reports [16]. These visualizations include performance indicators of the classification models, distributions of detected anomalies, and XAI-based explanations that highlight the factors influencing specific decisions [2, 5, 16].

Such capabilities support both immediate incident response and long-term strategic analysis of security trends within the CI/CD environment [6, 7]. By combining automated detection with explainable and auditable outputs, the system aims to enhance trust in AI-assisted security decisions and to facilitate their adoption in enterprise and regulated contexts [2, 5, 7].

5. Experimental Setup and Results Analysis

This section presents the experimental evaluation of the proposed AI-driven security architecture for CI/CD environments. The primary objective of the experiments is to validate the effectiveness of the selected methods and the proposed architectural integration under realistic operational conditions, including both normal operation and adversarial scenarios [3, 5, 12].

A. Experimental Environment

The experimental setup was deployed in a hybrid environment that combines synthetic workloads with data derived from realistic CI/CD workflows [12, 14, 15]. The test infrastructure includes a version control system, a CI/CD orchestrator, container-based execution environments, and monitoring components, reflecting a typical DevSecOps pipeline architecture [12, 13, 15, 16].

The analytical components were executed in an isolated environment to avoid interference with operational processes and to ensure reproducibility of results, in line with recommended security and auditing practices [6, 7]. Data collection was performed both in near real time and from historical logs, enabling evaluation of the system in both online and offline analysis modes [12].

B. Dataset Description

The dataset used in the experiments consists of several categories of data commonly observed in CI/CD environments. These include version control data such as commit histories and file change statistics, CI/CD pipeline logs covering build, test, and deployment stages, and Infrastructure as Code (IaC) configurations with both valid and intentionally vulnerable parameter settings [12, 13]. In addition, container metadata and dependency information were collected to reflect realistic software supply chain conditions [12, 15].

To evaluate robustness and detection capability, the dataset also includes synthetically generated threat scenarios, such as simulated supply chain attacks, adversarial inputs targeting the ML models, and examples of AI-assisted social engineering patterns [3, 5, 11]. All data were preprocessed using normalization and feature extraction procedures consistent with established ML security analytics practices [3, 12].

C. Experimental Scenarios

The evaluation was conducted using a set of representative operational and adversarial scenarios. These scenarios include normal pipeline execution with expected behavior, configuration errors introduced into IaC scripts, simulated supply chain compromises through modified dependencies, and adversarial inputs designed to test model robustness [3, 5, 9, 12]. In addition, social engineering-related scenarios were included to assess the system's ability to detect patterns associated with AI-assisted phishing and identity impersonation attempts targeting DevOps personnel [11, 12].

This scenario-based approach ensures that the evaluation reflects both routine operational conditions and high-risk threat situations commonly discussed in the literature on CI/CD and DevSecOps security [3, 9, 12].

D. Evaluation Metrics

The performance of the detection and classification components was assessed using standard metrics commonly applied in security-related machine learning studies, including precision, recall, F1-score, and ROC-based measures [3, 12]. For anomaly detection, reconstruction error distributions and threshold-based detection rates were analyzed to distinguish normal from abnormal behavior [12].

In addition to predictive performance, the quality of explanations generated by the XAI components was evaluated qualitatively, focusing on whether the highlighted features were consistent with domain knowledge and operational expectations [2, 5]. The impact of federated

learning on model generalization and data privacy was assessed by comparing centralized and distributed training setups [4, 7, 8].

E. Results and Discussion

The experimental results indicate that the combined use of supervised and unsupervised models provides robust detection capabilities across a wide range of scenarios, including previously unseen anomalies [3, 12]. In particular, autoencoder-based anomaly detection proved effective in identifying configuration errors and unusual pipeline behaviors, while supervised classifiers achieved reliable performance in categorizing known threat types [3, 12].

The integration of Explainable AI methods, specifically LIME and SHAP, significantly improved the interpretability of model decisions, enabling analysts to trace detections back to meaningful features in logs, configurations, and metadata [2, 5, 17, 18]. This property is essential for operational adoption and compliance with auditing requirements emphasized in security standards and regulations [6, 7].

Federated learning experiments demonstrated that collaborative training across multiple environments can preserve detection performance while reducing the need to centralize sensitive operational data, thus supporting privacy and governance requirements [4, 7, 8]. Finally, robustness testing under adversarial conditions confirmed that adversarial training and systematic stress testing improve model resilience, although no approach can guarantee complete immunity to sophisticated attacks [5, 6].

Overall, the results support the feasibility of the proposed architecture as a proactive, explainable, and privacy-aware security solution for modern CI/CD pipelines, consistent with the objectives outlined in DevSecOps-oriented security frameworks [12, 13].

6. Conclusion

This paper presented an integrated, AI-driven approach to enhancing the security of CI/CD pipelines in modern DevSecOps environments. The proposed framework combines supervised and unsupervised machine learning, Explainable Artificial Intelligence, federated learning, and adversarial robustness mechanisms into a unified architectural model designed to address the dynamic and complex threat landscape of automated software delivery systems.

The analysis demonstrates that traditional, rule-based security controls are increasingly insufficient for highly automated and rapidly evolving CI/CD infrastructures. In contrast, data-driven methods enable scalable anomaly detection and risk classification across heterogeneous data sources, including code repositories, pipeline logs, configuration files, and container metadata. The integration of Explainable AI techniques ensures that automated decisions remain transparent and auditable, which is essential for operational trust and regulatory compliance in security-critical environments.

Furthermore, the adoption of federated learning addresses key data governance and privacy constraints by enabling collaborative model training without centralizing sensitive operational data. This property is particularly relevant for multi-environment and multi-organizational CI/CD ecosystems subject to diverse regulatory requirements. The inclusion of adversarial robustness considerations highlights the necessity of treating machine learning models as potential attack surfaces and of continuously testing and adapting them against evolving threats.

The experimental evaluation indicates that the proposed architecture can effectively detect a broad range of security-relevant events, provide meaningful explanations for its decisions, and maintain acceptable performance under privacy-preserving and adversarial conditions. Although no technical solution can guarantee complete protection, the results support the conclusion that AI-assisted, explainable, and privacy-aware security analytics represent a practical and forward-looking direction for strengthening CI/CD and DevSecOps security.

Future work will focus on large-scale industrial validation, deeper integration with organizational governance processes, and the exploration of advanced model lifecycle management techniques to further improve robustness, transparency, and operational reliability.

REFERENCES

1. C. Cath *et al.*, “Artificial Intelligence and the ‘Good Society’: The US, EU, and UK Approach,” *Science and Engineering Ethics*, vol. 24, no. 2, pp. 505–528, 2018.
2. D. Gunning *et al.*, “XAI—Explainable Artificial Intelligence,” *Science Robotics*, vol. 4, no. 37, 2019.
3. D. Ucci, L. Aniello, and R. Baldoni, “Survey of Machine Learning Techniques for Malware Analysis,” *Computers & Security*, vol. 81, pp. 123–147, 2019.
4. H. Yang *et al.*, “Federated Machine Learning: Concept and Applications,” *ACM Transactions on Intelligent Systems and Technology*, vol. 10, no. 2, pp. 1–19, 2019.
5. F. Doshi-Velez and B. Kim, “Towards a Rigorous Science of Interpretable Machine Learning,” *arXiv preprint arXiv:1702.08608*, 2017.
6. National Institute of Standards and Technology (NIST), *Security and Privacy Controls for Information Systems and Organizations*, NIST SP 800-53 Rev. 5, 2020.
7. ISO/IEC, *ISO/IEC 27001:2022 Information Security, Cybersecurity and Privacy Protection — Information Security Management Systems — Requirements*, 2022.
8. European Union, *Digital Operational Resilience Act (DORA)*, Regulation (EU) 2022/2554, 2022.
9. I. Kartvelishvili and G. Kuchava, “Optimization of Software Delivery Quality and Speed in DevOps Using CI/CD,” in *Proc. Georgian Technical University Conference*, Tbilisi, 2024.
10. I. Kartvelishvili, M. Okhanashvili, and N. Chorkhauri, “Review and Analysis of Existing Methods for Network Attack Detection,” in *Proc. International Scientific-Practical Conference*, Tbilisi, 2023.
11. A. Bichnigauri *et al.*, “Strengthening Cyber Defenses — The Crucial Role of Phishing Simulation in Modern Security Strategies,” *Defence and Science*, no. 3, 2024, doi: 10.61446/ds.3.2024.8467.
12. Google Cloud, “MLOps: Continuous Delivery and Automation Pipelines in Machine Learning,” Google Cloud Documentation, 2022.
13. Microsoft, “Secure DevOps Kit for Azure (AzSK),” Microsoft Documentation, 2023.
14. TensorFlow, “TensorFlow Extended (TFX) Documentation.” [Online]. Available: <https://www.tensorflow.org/tfx>
15. Kubeflow, “Kubeflow Pipelines Documentation.” [Online]. Available: <https://www.kubeflow.org/docs/components/pipelines>
16. Grafana Labs, “Grafana Documentation.” [Online]. Available: <https://grafana.com/docs/>

17. M. T. Ribeiro, S. Singh, and C. Guestrin, “LIME: Local Interpretable Model-Agnostic Explanations.” [Online]. Available: <https://github.com/marcotcr/lime>
18. S. Lundberg and S.-I. Lee, “SHAP: A Unified Approach to Interpreting Model Predictions.” [Online]. Available: <https://github.com/slundberg/shap>
19. European Union, “Artificial Intelligence Act (EU AI Act).” [Online]. Available: <https://artificialintelligenceact.eu/>