Frontiers in Emerging Computer Science and Information Technology (Open Access)

Volume 03, Issue 08, August 2026,
Publish Date: 26-08-2026
DOI: https://doi.org/10.64917/fecsit/Volume03Issue08-01 **PageNo.06-12**

A Hybrid Deep Learning Framework for Real-Time Anomaly and Zero-Day Attack Detection in Cyber-Physical Manufacturing Systems

Faisal Al-Harbi
Department of Cybersecurity and Intelligent Systems, Advanced Technology University, Riyadh, Saudi Arabia

Nora Al-Qahtani
Department of Artificial Intelligence and Security Analytics, Institute of Digital Security Studies, Jeddah, Saudi Arabia.

ABSTRACT

Cyber-physical manufacturing systems (CPMSs) integrate industrial control, networked communication, sensing, computation, and automated decision-making, creating highly interconnected production environments but also expanding the attack surface available to sophisticated cyber threats. Conventional intrusion detection approaches frequently depend on predefined signatures or static classification boundaries, making them less effective against previously unseen and rapidly evolving attacks. This paper proposes a hybrid deep learning framework for real-time anomaly and zero-day attack detection in CPMSs. The proposed architecture combines heterogeneous feature extraction, temporal representation learning, anomaly scoring, ensemble decision fusion, and adaptive classification to distinguish normal operational deviations from malicious behavior. Its theoretical foundation integrates sequence-based deep learning with ensemble decision mechanisms and imbalance-aware learning. Insights from existing work on malicious URL detection, intrusion detection, traffic classification, deep learning, decision fusion, and oversampling are synthesized to formulate a CPMS-oriented architecture. The framework emphasizes low-latency inference, resilience to unknown attack patterns, and operational interpretability. Rather than claiming experimentally measured performance without an executed benchmark, the Results section presents analytically derived findings and expected evaluation behavior. The proposed framework provides a research foundation for developing adaptive security mechanisms capable of detecting both known and zero-day threats while maintaining compatibility with real-time manufacturing constraints.

KEYWORDS: Cyber-Physical Manufacturing Systems, Zero-Day Attack Detection, Anomaly Detection, Deep Learning, Intrusion Detection, Industrial IoT, Ensemble Learning, Real-Time Security, Adaptive Learning.

INTRODUCTION

The convergence of information technology and operational technology has transformed manufacturing systems from isolated industrial environments into interconnected cyber-physical ecosystems. Modern manufacturing infrastructures increasingly depend on networked sensors, programmable controllers, industrial computers, communication gateways, software-defined networking, and intelligent analytics. This connectivity improves automation and operational efficiency but simultaneously introduces security dependencies across multiple computational and communication layers. An intrusion affecting a network component may consequently influence physical processes, production continuity, or automated control decisions.

A fundamental difficulty in protecting such environments is the distinction between known malicious behavior and previously unseen behavior. Signature-oriented security mechanisms are effective when an attack has already been characterized, but they are intrinsically constrained when confronted with novel attack structures. Machine learning and deep learning have therefore become important approaches for identifying complex patterns in network and application data. Liu and Lang's survey demonstrates the increasing relevance of machine learning and deep learning techniques for intrusion detection systems, while also highlighting the diversity of learning strategies and detection architectures available for security applications (Liu and Lang, 2019).

The problem becomes more challenging in CPMSs because normal system behavior is itself highly dynamic. A temporary increase in network traffic, unusual sensor measurements, maintenance activity, or changes in production schedules may resemble malicious behavior without representing an actual attack. Consequently, a practical detector must identify deviations relative to operational context rather than treating every unusual observation as malicious. The distinction between statistical anomaly and security-relevant anomaly is therefore central to the design of an effective detection framework.

Research on malicious URL detection provides useful methodological evidence for this problem. Aljabri et al. investigated lexical, network, and content-based characteristics using machine learning and deep learning, demonstrating the value of combining heterogeneous feature representations rather than relying on a single information source (Aljabri et al., 2022). Their work is not directly a CPMS study, but its feature-fusion principle can be conceptually transferred to industrial security, where network packets, protocol attributes, temporal behavior, sensor observations, and device-level information can provide complementary evidence.

Similarly, decision-fusion approaches can strengthen security classification when individual models possess different strengths. Alanazi and Gumaei proposed an ensemble decision-fusion approach for malicious website detection, providing a basis for considering multiple predictive perspectives rather than relying exclusively on a single classifier (Alanazi and Gumaei, 2023). This principle is particularly relevant to CPMS security because attacks may simultaneously exhibit temporal, statistical, protocol-level, and behavioral characteristics.

Recent research also emphasizes the importance of security analysis across heterogeneous computing environments. Govindarajan et al. presented SecureRiskNet as an AI-driven framework for intelligent security-risk detection across heterogeneous cloud-fog computing networks, reinforcing the broader principle that intelligent security mechanisms should account for distributed and heterogeneous infrastructures rather than treating security events as isolated observations (Govindarajan et al., 2026). Although cloud-fog environments differ from manufacturing systems, their heterogeneous and distributed characteristics provide a useful theoretical basis for designing security analytics capable of operating across multiple computational layers.

The objective of this paper is therefore to formulate a hybrid deep learning framework that combines anomaly detection and zero-day attack identification within a unified CPMS security architecture. The framework is designed around five requirements: heterogeneous feature integration, temporal behavior modeling, adaptive anomaly detection, decision-level fusion, and real-time inference. The paper contributes a structured conceptual methodology, a technical architecture, an evaluation framework, and a critical analysis of the limitations associated with deep learning-based security detection.

LITERATURE REVIEW

Existing research demonstrates that security detection performance depends substantially on feature representation, model diversity, data quality, and the ability to identify nonlinear behavioral patterns. Aljabri et al. examined lexical, network, and content-based characteristics for malicious URL detection, showing that security classification can benefit from combining heterogeneous evidence (Aljabri et al., 2022). In a CPMS, an analogous strategy would combine network-flow attributes, industrial protocol characteristics, device states, temporal sequences, and process-level observations.

Aljabri et al. further reviewed machine learning techniques for malicious URL detection and identified the continuing importance of feature engineering and learning-model selection in security applications (Aljabri et al., 2022). The relevance to manufacturing security lies not in the specific URL domain but in the underlying detection problem: malicious behavior may be expressed through multiple observable characteristics, and a robust detection system should avoid depending on one narrow feature representation.

Yang et al. developed a convolutional gated-recurrent-unit neural network for malicious URL detection (Yang et al., 2019). The combination of convolutional and recurrent operations is theoretically important for CPMS detection. Convolutional components can identify local patterns within structured observations, while recurrent mechanisms can model temporal dependencies. Manufacturing attacks frequently evolve over sequences rather than isolated events; therefore, temporal modeling is particularly important.

Lee et al. investigated optimization and machine learning classifiers for malicious URL detection, further supporting the use of algorithmic classification for identifying complex malicious patterns (Lee et al., 2020). Mourtaji et al. similarly explored a hybrid rule-based and convolutional neural network solution for phishing URL detection (Mourtaji et al., 2021). These studies indicate that purely statistical learning does not necessarily have to replace deterministic security rules. Instead, hybrid architectures can combine learned patterns with explicit operational knowledge.

Decision fusion represents another important research direction. Alanazi and Gumaei demonstrated an ensemble

decision-fusion approach for malicious website detection (Alanazi and Gumaei, 2023). For CPMS environments, this concept can be extended by allowing multiple models to independently evaluate network anomalies, temporal deviations, and process-level behavior before producing a final risk decision.

Data imbalance presents an additional challenge. Security datasets generally contain substantially fewer attack samples than normal observations, and this problem can become more severe for zero-day attacks because genuinely novel attack instances may be extremely scarce. Le et al. investigated oversampling techniques for prediction under imbalanced conditions, providing methodological support for carefully addressing minority-class representation (Le et al., 2018). However, synthetic oversampling must be applied cautiously in security applications because artificially generated observations may not faithfully represent realistic attack trajectories.

The broader methodological foundation is supported by Liu and Lang's analysis of machine learning and deep learning methods for intrusion detection systems (Liu and Lang, 2019). Their work establishes the relevance of learning-based approaches while also emphasizing the diversity of detection methods and the continuing challenge of selecting appropriate models for intrusion detection.

Some provided studies address domains outside cybersecurity, including plant disease classification and adaptive neuro-fuzzy inference. Anagnostis et al. demonstrated the use of deep learning for image-based classification, while Bayraktar et al. investigated adaptive neuro-fuzzy inference for disease detection (Anagnostis et al., 2021; Bayraktar et al., 2021). Although these applications are not directly transferable to CPMS intrusion detection, they support the broader theoretical observation that nonlinear learning and adaptive inference can classify complex patterns.

Mohammed et al. investigated machine learning and deep learning for traffic classification and prediction in software-defined networking (Mohammed et al., 2019). This is particularly relevant to CPMSs because industrial networks require continuous traffic characterization and abnormal-flow identification.

Collectively, the literature reveals a gap between general machine-learning intrusion detection and an integrated CPMS architecture capable of simultaneously addressing temporal anomalies, heterogeneous observations, class imbalance, decision uncertainty, and zero-day behavior. The proposed framework addresses this gap through hybrid representation learning and decision fusion.

METHODOLOGY

3.1 Framework Architecture

The proposed framework consists of six logical layers: data acquisition, preprocessing, heterogeneous feature extraction, hybrid deep representation learning, anomaly and classification engines, and decision fusion.

At the data acquisition layer, observations are collected from industrial network flows, communication gateways, sensors, controllers, endpoint devices, and manufacturing applications. Each observation is associated with a timestamp so that sequential dependencies can be preserved. The objective is not merely to collect packet-level information but to construct a multidimensional representation of system behavior.

The preprocessing layer normalizes numerical attributes, encodes categorical variables, removes corrupted observations, and organizes data into temporal windows. Let an observation at time tt be represented as

Xt=[Nt,Pt,St,Dt],X_t=[N_t,P_t,S_t,D_t],

where NtN_t represents network attributes, PtP_t represents protocol-level characteristics, StS_t represents system or sensor states, and DtD_t represents device-related attributes.

A sequence of observations is consequently represented as

Xt−k:t={Xt−k,...,Xt−1,Xt}.\mathbf{X}_{t-k:t}=\{X_{t-k},...,X_{t-1},X_t\}.

This representation enables the framework to evaluate both instantaneous and sequential abnormalities.

3.2 Heterogeneous Feature Extraction

The feature-extraction layer generates complementary representations from different data modalities. Network features may include packet rates, connection durations, byte distributions, protocol frequencies, and communication relationships. Temporal features characterize changes in these variables over consecutive windows. Device and process features capture deviations from expected operational states.

The rationale for heterogeneous extraction follows the feature-combination principle demonstrated in malicious URL detection research, where lexical, network, and content features provide complementary information (Aljabri et al., 2022). In a manufacturing environment, the equivalent principle is to avoid assuming that an attack will be visible through only network traffic or only process measurements.

3.3 Hybrid Deep Representation Learning

The central learning component combines convolutional and recurrent representations. A convolutional neural

network (CNN) identifies local relationships among features, while a gated recurrent unit (GRU) captures temporal dependencies.

Conceptually,

HtCNN=fCNN(Xt)H_t^{CNN}=f_{CNN}(X_t)

and

HtGRU=fGRU(Ht−k:tCNN).H_t^{GRU}=f_{GRU}(H_{t-k:t}^{CNN}).

The resulting representation is

Ht=[HtCNN‖HtGRU],H_t=[H_t^{CNN}\parallel H_t^{GRU}],

where ‖\parallel denotes feature concatenation.

The CNN-GRU structure is motivated by the architecture investigated by Yang et al. for malicious URL detection (Yang et al., 2019). Its adaptation to CPMSs is theoretically appropriate because manufacturing attacks can contain both localized signatures and temporal progression.

For example, a sudden protocol deviation may initially appear as a local statistical anomaly. If that deviation is followed by repeated unauthorized communication and abnormal controller activity, the sequence becomes substantially more suspicious. A temporal model can distinguish this progression from a single legitimate operational event.

3.4 Dual Detection Mechanism

The framework uses two complementary detection mechanisms.

The first is a supervised classification branch that learns known attack categories. Let

P(y=c|Ht)P(y=c|H_t)

denote the probability that representation HtH_t belongs to attack class cc. This branch is optimized using labeled observations.

The second is an anomaly-detection branch designed to identify observations that do not conform to learned normal behavior. An autoencoder-based formulation can reconstruct normal representations:

X^t=D(E(Xt)),\hat{X}_t=D(E(X_t)),

where EE is the encoder and DD is the decoder. The reconstruction error is

At=‖Xt−X^t‖2.A_t=\|X_t-\hat{X}_t\|^2.

Large values of AtA_t indicate that the observation differs substantially from the learned normal distribution.

This dual strategy is essential for zero-day detection. A supervised classifier may fail when an attack does not belong to any known class, whereas an anomaly detector can flag previously unseen behavior. Conversely, anomaly detection alone may generate excessive false positives because legitimate manufacturing changes can also appear unusual. Combining both mechanisms therefore provides a more balanced security strategy.

3.5 Adaptive Decision Fusion

The final detection decision is generated by combining classifier confidence and anomaly severity:

Rt=αCt+(1−α)At,R_t=\alpha C_t+(1-\alpha)A_t,

where CtC_t represents supervised classification confidence, AtA_t represents normalized anomaly severity, and α\alpha controls their relative contribution.

The framework can further introduce a temporal consistency factor:

Rt′=βRt+(1−β)Rt−1.R_t'=\beta R_t+(1-\beta)R_{t-1}.

This prevents a single isolated observation from immediately producing a high-confidence security decision unless its severity exceeds a predefined threshold.

Decision fusion is theoretically supported by the ensemble approach of Alanazi and Gumaei (2023). In heterogeneous distributed infrastructures, intelligent security decisions similarly benefit from combining multiple evidence sources rather than relying on one model. Govindarajan et al. (2026) provide a related conceptual basis through AI-driven security-risk detection across heterogeneous cloud-fog environments, where distributed infrastructure requires intelligent risk assessment across multiple system components.

3.6 Handling Class Imbalance

Because normal manufacturing traffic may dominate attack traffic, training must account for class imbalance. Oversampling can be applied to minority attack categories, with appropriate restrictions against generating unrealistic attack sequences. The methodology is informed by Le et al. (2018), whose work demonstrates the importance of oversampling strategies in imbalanced prediction problems.

A cost-sensitive objective can additionally be formulated as

L=−∑c=1Cwcyclog⁡(y^c),L=-\sum_{c=1}^{C}w_c y_c\log(\hat{y}_c),

where wcw_c assigns higher importance to underrepresented attack classes.

For zero-day evaluation, synthetic oversampling should not substitute for genuinely unseen attack testing. Otherwise, the evaluation could inadvertently measure the model's ability to recognize artificially generated variations of known attacks rather than true generalization to novel threats.

3.7 Real-Time Deployment

Real-time operation requires a balance between detection accuracy and computational cost. The framework therefore uses a lightweight feature-processing pipeline followed by optimized inference. A hierarchical deployment model can place preprocessing and rapid anomaly scoring at the edge while forwarding higher-risk observations to more computationally intensive analysis.

This architecture is consistent with the broader distributed-security perspective discussed in heterogeneous cloud-fog security research (Govindarajan et al., 2026). In CPMSs, such distribution can reduce response latency and limit the amount of raw industrial data transferred across the network.

3.8 Evaluation Strategy

The proposed framework should be evaluated using accuracy, precision, recall, F1-score, false-positive rate, detection latency, and computational overhead. Accuracy alone is insufficient because highly imbalanced security datasets can produce deceptively high accuracy even when minority attacks are poorly detected.

The most important experiment should be a zero-day split, where selected attack families are completely excluded from training and validation. The detector is then evaluated on those unseen families. This design directly tests generalization rather than conventional memorization.

Ablation studies should separately remove the CNN component, GRU component, anomaly branch, decision-fusion layer, and imbalance-handling mechanism. Such experiments would identify which architectural components contribute meaningfully to detection capability.

RESULTS

Because the framework presented here is a proposed architecture rather than a report of an executed experimental campaign, the findings are analytical rather than fabricated numerical results. The architecture indicates several expected performance characteristics that should be empirically validated in future implementation.

First, the hybrid CNN-GRU structure should provide stronger temporal discrimination than a purely static classifier. CNN-based feature extraction can capture localized correlations, while GRU-based processing can recognize sequential dependencies. This is particularly important when an attack evolves through multiple stages instead of producing one immediately recognizable event. The relevance of combining convolutional and recurrent processing is supported by the architecture explored by Yang et al. (2019).

Second, the dual supervised-anomaly structure is expected to improve resilience against unknown attacks. A classifier provides strong discrimination for known attack classes, whereas reconstruction-based anomaly scoring provides an additional mechanism for identifying behavior outside the learned normal distribution. The combined approach should therefore reduce dependence on predefined attack labels.

Third, decision fusion is expected to improve robustness when different indicators provide conflicting evidence. For example, abnormal network activity without corresponding process disruption may represent legitimate maintenance, whereas simultaneous network deviation and unusual device behavior provides stronger evidence of compromise. Ensemble decision principles studied by Alanazi and Gumaei (2023) support this type of multi-evidence reasoning.

Fourth, class-imbalance handling should improve sensitivity to rare attacks, although excessive oversampling may increase false-positive behavior if synthetic observations do not accurately represent operational conditions. Consequently, imbalance mitigation should be evaluated jointly with precision and false-positive rate rather than recall alone.

Finally, real-time deployment remains dependent on model complexity, feature volume, hardware resources, and communication latency. The proposed edge-oriented architecture can reduce response time by performing initial screening close to the industrial process. This distributed strategy is conceptually consistent with intelligent risk detection in heterogeneous computing infrastructures described by Govindarajan et al. (2026).

These findings should not be interpreted as measured improvements over existing methods. Instead, they establish testable hypotheses for subsequent experimental validation using representative CPMS datasets and controlled zero-day attack scenarios.

DISCUSSION

The proposed framework addresses a fundamental weakness of conventional intrusion detection: the assumption that malicious behavior can be adequately represented by previously observed attack classes. In CPMSs, this assumption is particularly problematic because

attackers can alter payloads, communication patterns, timing behavior, or attack sequences while retaining the same underlying objective. A model that learns only fixed attack signatures may therefore achieve strong conventional classification performance while remaining vulnerable to previously unseen behaviors.

The hybrid architecture provides a theoretical response to this problem by separating recognition from deviation detection. Supervised learning answers the question, "Does this behavior resemble a known attack?" Anomaly detection addresses a different question, "Does this behavior differ substantially from expected operational behavior?" These questions are complementary rather than interchangeable.

The literature also suggests that model diversity is valuable. Alanazi and Gumaei (2023) demonstrated decision fusion in malicious website detection, while Liu and Lang (2019) documented the broad applicability of machine learning and deep learning to intrusion detection. The proposed CPMS framework extends these principles by combining feature-level, temporal, and decision-level representations.

Nevertheless, anomaly detection introduces a critical trade-off. Manufacturing environments are dynamic by design. Production schedules, machine configurations, maintenance operations, and workload conditions may change legitimately. A detector trained on a narrow definition of normality could therefore classify legitimate operational changes as attacks. Adaptive learning must consequently be constrained by operational safeguards; unrestricted online model updates could allow an attacker to gradually manipulate the learned baseline.

Zero-day detection presents an even greater evaluation challenge. A system cannot credibly demonstrate zero-day capability merely by randomly dividing a dataset into training and testing partitions. If closely related instances of the same attack appear in both sets, the model may effectively recognize familiar patterns. Attack-family exclusion and temporal separation are therefore necessary for meaningful evaluation.

Interpretability is another important practical consideration. A high anomaly score alone may not be sufficient for an industrial operator to initiate a production-impacting response. The system should identify contributing evidence, such as unusual communication frequency, protocol deviation, abnormal temporal behavior, or inconsistent device states. This requirement is especially important in safety-sensitive manufacturing environments.

The framework also has limitations. Deep models may require substantial training data, computational resources, and careful hyperparameter selection. Rare zero-day attacks cannot easily be represented during training, making unsupervised or semi-supervised components necessary but potentially less precise. Furthermore, the references available for this study include substantial evidence from adjacent security domains rather than CPMS-specific experimental datasets. Therefore, direct transfer of reported methodologies should be regarded as theoretical adaptation rather than empirical equivalence.

A further limitation is that cybersecurity performance cannot be reduced to a single metric. Increasing recall may increase false alarms, while aggressive latency optimization may reduce analytical depth. The most appropriate deployment objective is therefore a multi-dimensional balance among detection capability, false-positive behavior, latency, computational cost, and operational safety. Govindarajan et al. (2026) similarly reinforce the importance of intelligent security-risk analysis across heterogeneous infrastructures rather than treating security as a single isolated classification problem.

CONCLUSION

This paper proposed a hybrid deep learning framework for real-time anomaly and zero-day attack detection in cyber-physical manufacturing systems. The architecture integrates heterogeneous feature extraction, CNN-based local representation learning, GRU-based temporal modeling, anomaly detection, supervised attack classification, imbalance-aware training, and adaptive decision fusion.

The principal contribution is the integration of complementary detection perspectives within a single CPMS-oriented architecture. Known attacks can be recognized through supervised learning, while unknown behavior can be identified through anomaly scoring. Decision fusion provides a mechanism for combining these signals, and temporal modeling allows the system to distinguish isolated operational variations from persistent or evolving malicious behavior.

The literature indicates that heterogeneous feature representation, deep learning, ensemble decision-making, traffic analysis, and imbalance-aware learning each contribute useful principles to security analytics. However, these approaches must be carefully adapted to manufacturing conditions because legitimate operational variability creates a particularly difficult false-positive problem.

Future research should implement the proposed architecture using realistic industrial datasets and controlled cyber-physical testbeds. Evaluation should include strict zero-day attack-family separation, temporal validation, ablation studies, latency measurement, computational-resource analysis, and explainability

assessment. Future versions could additionally investigate federated learning, continual learning, graph-based communication modeling, and adaptive edge-cloud deployment. Such extensions could improve the framework's ability to operate across heterogeneous manufacturing environments while reducing dependence on centralized security analytics.

Ultimately, effective CPMS security requires more than accurate classification. It requires a detection architecture capable of recognizing known threats, identifying unfamiliar behavior, adapting to legitimate operational changes, and producing sufficiently timely and interpretable evidence for secure industrial decision-making.

REFERENCES

1. A. Alanazi and A. Gumaei, "A decision-fusion-based ensemble approach for malicious websites detection," Appl. Sci., vol. 13, no. 18, Sep. 2023, Art. no. 10260.
2. A. Anagnostis et al., "A deep learning approach for anthracnose infected trees classification in walnut orchards," Comput. Electron. Agric., vol. 182, Mar. 2021, Art. no. 105998.
3. M. M. Aljabri et al., "An assessment of lexical, network, and content-based features for detecting malicious urls using machine learning and deep learning models," Comput. Intell. Neurosci., vol. 2022, no. 1, 2022, Art. no. 3241216.
4. M. Aljabri et al., "Detecting malicious URLs using machine learning techniques: Review and research directions," IEEE Access, vol. 10, pp. 121395–121417, Aug. 2022.
5. R. Bayraktar, B. Haznedar, K. S. Bayram, and M. F. Hasoˇglu, "Plant disease detection by using adaptive neuro-fuzzy inference system," Tamap J. Eng., vol. 2021, no. 125, pp. 1–10, Sep. 2021.
6. N. Bhadouria, "Malicious_URL's_Dataset," 2022. Accessed: Jun. 5, 2023. [Online]. Available: https://www.kaggle.com/datasets/naveenbhadouria/malicious
7. Govindarajan, V., Ahmed, F., Kamaluddin, K. et al. SecureRiskNet: An Advanced AI-Driven Framework for Intelligent Security Risk Detection in Heterogeneous Cloud-Fog Computing Networks. Int J Comput Intell Syst 19, 74 (2026).
8. T. Le, M. Y. Lee, J. R. Park, and S. W. Baik, "Oversampling techniques for bankruptcy prediction: Novel features from a transaction dataset," Symmetry, vol. 10, no. 4, Mar. 2018, Art. no. 79.
9. O. V. Lee et al., "A malicious URLs detection system using optimization and machine learning classifiers," Indones J. Electr. Eng. Comput. Sci., vol. 17, no. 3, pp. 1210–1214, Mar. 2020.
10. H. Liu and B. Lang, "Machine learning and deep learning methods for intrusion detection systems: A survey," Appl. Sci., vol. 9, no. 20, Oct. 2019, Art. no. 4396.
11. A. R. Mohammed, S. A. Mohammed, and S. Shirmohammadi, "Machine learning and deep learning based traffic classification and prediction in software defined networking," in IEEE Int. Symp. Meas. Netw. (M&N), Catania, Italy, Jul. 8–10, 2019, pp. 1–6.
12. Y. Mourtaji, M. Bouhorma, D. Alghazzawi, G. Aldabbagh, and A. Alghamdi, "Hybrid rule-based solution for phishing URL detection using convolutional neural network," Wirel. Commun. Mob. Comput., vol. 2021, no. 1, pp. 1–24, 2021, Art. no. 8241104.
13. W. Yang, W. Zuo, and B. Cui, "Detecting malicious URLs via a keyword-based convolutional gated-recurrent-unit neural network," IEEE Access, vol. 7, pp. 29891–29900, Jan. 2019.
14. Ramamurthy, K., Gumber, S., Abdelfattah, W.M. et al. Human AI trust modeling in cognitive systems via ensemble learning and advanced feature engineering. Discov Artif Intell 6, 366 (2026). https://doi.org/10.1007/s44163-026-01255-7
15. Philip, P. G. (2025). Explainable Artificial Intelligence (XAI) for Project Governance: Improving Transparency and Stakeholder Trust in Automated Project Decision. Journal of Project Management Studies, 2(1), 37–55. https://doi.org/10.58425/jpms.v2i1.570