INTERNATIONAL JOURNAL OF DATA SCIENCE AND MACHINE LEARNING (ISSN: 2692-5141)
VOL.06 ISSUE 02 (2026)

Research Article

Adaptive Ensemble-Based Intrusion Detection with Explainable AI for Secure Drone and Industrial Cyber-Physical Systems

Dr. Ishara Fernando [1]

[1]Department of Computer Science, Kandy Institute of Digital Technology, Kandy, Sri Lanka

Received: 25 May 2026
Revised: 20 June 2026
Accepted: 14 July 2026
Published: 27 August 2026

Doi:10.55640/ijdsml-06-02-06

Page No: 140-148

Abstract

The convergence of unmanned aerial systems, industrial control infrastructure, communication networks, and cyber-physical processes has created heterogeneous environments in which intrusion detection must operate across diverse traffic characteristics, computational constraints, and attack classes. Conventional intrusion detection approaches may struggle to maintain reliable discrimination when network conditions, device behavior, and attack distributions change. This paper proposes an adaptive ensemble-based intrusion detection framework integrating heterogeneous detection models, evolutionary optimization, explainable artificial intelligence (XAI), and context-sensitive decision fusion for drone and industrial cyber-physical systems. The theoretical foundation combines ensemble learning with genetic-algorithm-based optimization, motivated by established work on genetic algorithms and optimization-oriented communication-system design (Katoch, Chauhan and Kumar, 2021; Kramer and Kramer, 2017). The framework further incorporates adaptive model selection and explainability so that security operators can distinguish benign operational deviations from malicious activity. The directly related work on explainable ensemble intrusion detection in heterogeneous drone and industrial networks provides an important positioning reference for this research (Islam, Ahmed, Ishtiaq et al., 2026). The proposed methodology emphasizes multi-class detection, feature-level diversity, adaptive weighting, evolutionary optimization, and explanation consistency. Analytical findings indicate that the combination of complementary classifiers and adaptive weighting can theoretically address weaknesses associated with single-model detection, while XAI can improve operational interpretability. However, the framework also introduces computational, calibration, and explanation-consistency challenges. The study therefore positions adaptive ensemble detection as a promising architecture for security-sensitive cyber-physical environments where detection accuracy must be balanced against latency, computational cost, and human interpretability.

Keywords: Adaptive Ensemble Learning; Intrusion Detection; Explainable Artificial Intelligence; Drone Networks; Industrial Cyber-Physical Systems; Genetic Algorithms; Multi-Class Classification; Cybersecurity; Adaptive Weighting; Network Security

INTRODUCTION

The increasing integration of communication networks with physical processes has transformed both drone systems and industrial environments into complex cyber-physical ecosystems. In such environments, security events are not restricted to conventional information-technology networks. A malicious communication pattern can potentially affect control decisions, operational continuity, resource allocation, or

physical processes. The security challenge is consequently twofold: intrusion detection must identify abnormal network behavior, while the resulting decision must be sufficiently interpretable to support operational intervention.

Heterogeneity is a central difficulty. Drone networks may involve mobile nodes, variable communication conditions, resource-constrained devices, and rapidly changing traffic patterns, whereas industrial cyber-physical systems may contain relatively stable control processes alongside highly specialized communication protocols. A detector optimized for one operational regime may therefore become less effective when exposed to another. This problem motivates adaptive architectures in which multiple detection mechanisms contribute to a final decision rather than relying exclusively on one classifier.

The use of optimization techniques provides one possible foundation for this adaptation. Genetic algorithms have been extensively studied as general optimization mechanisms and can be employed for searching large solution spaces in which deterministic selection is difficult (Katoch, Chauhan and Kumar, 2021; Kramer and Kramer, 2017). Their relevance to communication-system design is also demonstrated by research applying genetic optimization to LDPC and polar-code construction (Elkelesh, Ebada, Cammerer and Ten Brink, 2019; Elkelesh, Ebada, Cammerer, Schmalen and Ten Brink, 2019). These studies do not directly solve intrusion detection, but they provide a theoretical basis for evolutionary optimization of model configurations, feature subsets, and ensemble weights.

The present research proposes an adaptive ensemble-based intrusion detection architecture designed specifically for heterogeneous drone and industrial cyber-physical environments. The framework combines complementary classifiers, adaptive weighting, evolutionary optimization, and explainability. Its objective is not merely to maximize classification performance but to construct a detection mechanism capable of adapting to changing operational contexts while providing interpretable evidence for its decisions.

The significance of this approach is reinforced by research specifically addressing explainable ensemble models for multi-class intrusion detection in heterogeneous drone and industrial networks (Islam, Ahmed, Ishtiaq et al., 2026). The present framework extends this research direction conceptually by emphasizing adaptive ensemble configuration and optimization as mechanisms for dealing with changing system conditions. The principal research objectives are therefore to establish an adaptive ensemble architecture, define an optimization-driven decision mechanism, incorporate explainability into the detection pipeline, and critically examine the trade-off between detection effectiveness, computational overhead, adaptability, and interpretability.

LITERATURE REVIEW

2.1 Optimization and Genetic-Algorithm Foundations

Genetic algorithms provide a population-based optimization mechanism capable of searching large and nonlinear solution spaces. Katoch, Chauhan and Kumar (2021) describe the development and broad applicability of genetic algorithms, while Kramer and Kramer (2017) provide a foundational treatment of genetic algorithms as computational optimization techniques. Their relevance to an intrusion detection architecture lies in their ability to optimize parameters without requiring a simple analytical relationship between the objective and decision variables.

The application of genetic optimization to communication systems further demonstrates the feasibility of using evolutionary search for highly structured technical problems. Elkelesh, Ebada, Cammerer and Ten Brink (2019) investigate decoder-tailored polar-code design using a genetic algorithm, while Elkelesh, Ebada, Cammerer, Schmalen and

Ten Brink (2019) apply genetic optimization to LDPC code design. Ebada, Elkelesh, Cammerer and ten Brink (2018) similarly examine finite-length LDPC code design using scattered EXIT charts. Although these studies address coding rather than cybersecurity, they demonstrate a common methodological principle: evolutionary search can be used to identify configurations that improve performance under complex constraints.

For intrusion detection, this principle can be transferred to ensemble construction. Candidate solutions may represent classifier weights, feature subsets, model combinations, or threshold configurations. The optimization objective can then jointly consider detection quality and computational cost.

2.2 Ensemble and High-Throughput Detection Foundations

Ferraz, Subramaniyan, Chinthala, Andrade, Cavallaro, Nandy and Falcao (2021) examine high-throughput non-binary LDPC decoders and compare ASIC, FPGA, and GPU architectures. Their work highlights the importance of implementation architecture when sophisticated algorithms must operate under high-throughput requirements. This consideration is directly relevant to cyber-physical intrusion detection because an algorithm with strong theoretical discrimination may still be unsuitable when processing latency becomes excessive.

Shao, Hailes, Wang, Wu, Maunder, Al-Hashimi and Hanzo (2019) survey ASIC implementations for turbo, LDPC, and polar decoders and further demonstrate the relationship between algorithmic complexity and hardware realization. These studies motivate an important design principle for intrusion detection: the ensemble should not be optimized exclusively for classification quality. Its computational footprint must also be considered.

2.3 Cyber-Physical and Industrial Security Context

Eremia, Liu and Edris (2016) discuss advanced power-system solutions involving HVDC, FACTS, and artificial intelligence. Their work illustrates the increasing role of intelligent computational mechanisms in critical infrastructure and therefore reinforces the need for security mechanisms capable of operating alongside complex automated infrastructure.

Zhou and Shi (2024) specifically examine optimization of whitelisting technology for network firewalls in industrial control systems using genetic algorithms. This reference is particularly relevant because it connects evolutionary optimization with industrial network security. It suggests that genetic algorithms can support adaptive security policy optimization rather than being restricted to conventional engineering optimization.

Rao (2015) provides a broader foundation in channel-coding techniques for wireless communications, while Van Wonterghem, Alloum, Boutros and Moeneclaey (2016) analyze performance differences among short-length error-correcting codes. These works emphasize the importance of performance under constrained communication conditions, which is pertinent to drone systems in which communication resources may be variable or limited.

Wijekoon, Viterbo, Hong, Micheloni and Marelli (2019) investigate graph expansion and decoding for non-binary LDPC codes. Collectively, these communication-oriented studies reinforce the importance of robust algorithm design under constrained and heterogeneous communication environments.

2.4 Research Gap

The literature demonstrates strong foundations in evolutionary optimization, high-throughput algorithm implementation, industrial security optimization, and

communication-system robustness. However, the provided literature does not collectively establish a complete adaptive intrusion detection architecture that simultaneously integrates heterogeneous ensembles, evolutionary adaptation, multi-class intrusion discrimination, and explainable decision-making.

The most directly aligned study is Islam, Ahmed, Ishtiaq et al. (2026), which addresses advanced explainable ensemble models for multi-class intrusion detection in heterogeneous drone and industrial networks. This work establishes the relevance of explainable ensemble detection to the exact application context. The remaining references provide complementary methodological foundations rather than direct intrusion-detection solutions. The research gap therefore concerns the integration of adaptive optimization into explainable ensemble detection so that model composition and decision weights can respond to heterogeneous operating conditions (Islam, Ahmed, Ishtiaq et al., 2026).

METHODOLOGY

3.1 Proposed Adaptive Ensemble Architecture

The proposed architecture consists of five conceptual layers: data acquisition, feature representation, heterogeneous base learners, adaptive ensemble optimization, and explainable decision output.

At the first layer, telemetry and communication observations from drone and industrial cyber-physical environments are transformed into structured detection instances. These observations may represent communication behavior, protocol characteristics, temporal patterns, control-related variables, or other measurable network attributes. Because drone and industrial systems possess different operational characteristics, the framework treats the resulting observations as heterogeneous rather than assuming that one statistical distribution describes the complete environment.

The second layer performs feature representation. Instead of relying on a single feature configuration, the architecture permits multiple feature subsets or representations. This allows different classifiers to specialize in different behavioral dimensions. The principle is consistent with the broader optimization literature, where alternative configurations can be searched rather than manually fixed (Katoch, Chauhan and Kumar, 2021; Kramer and Kramer, 2017).

The third layer contains heterogeneous base learners. Rather than requiring all learners to produce identical internal representations, each model is allowed to capture different aspects of the traffic. One learner may be more sensitive to statistical deviations, another to nonlinear class boundaries, and another to temporal or contextual changes. The ensemble therefore derives its strength from complementary errors rather than simple model duplication.

3.2 Adaptive Weighting Mechanism

Let the base classifiers be represented as C1,C2,…,CMC_1,C_2,\ldots,C_M. For an input vector xx, each classifier produces a class probability Pm(y|x)P_m(y|x). The ensemble decision can be represented as:

Pens(y|x)=∑m=1MwmPm(y|x)P_{ens}(y|x)=\sum_{m=1}^{M} w_m P_m(y|x)

where wmw_m denotes the adaptive contribution of classifier mm, subject to:

wm≥0,∑m=1Mwm=1.w_m\geq0,\qquad \sum_{m=1}^{M}w_m=1.

The principal distinction from a static ensemble is that wmw_m is not permanently fixed. It can be adjusted according to validation performance, operational context, class-specific reliability, computational cost, or recent detection behavior.

For example, a classifier that performs strongly on industrial control traffic may receive a larger weight during an industrial operating regime, while another classifier may become more influential when the observed traffic exhibits characteristics associated with drone communication. Such adaptation reduces the assumption that a single classifier remains optimal under all circumstances.

3.3 Genetic Optimization of Ensemble Configuration

Genetic algorithms can be used to search for an effective ensemble configuration. A chromosome can encode classifier selection, feature-subset selection, model weights, and decision thresholds. The optimization objective may be expressed as:

F=αD−βL−γC+δXF=\alpha D-\beta L-\gamma C+\delta X

where DD represents detection effectiveness, LL represents false-alarm or misclassification loss, CC represents computational cost, and XX represents an explainability-related quality measure. The coefficients α,β,γ,δ\alpha,\beta,\gamma,\delta define the relative priorities.

The optimization process begins with a population of candidate configurations. Selection retains configurations with stronger fitness, crossover combines promising configurations, and mutation introduces diversity. This mechanism is theoretically consistent with established genetic-algorithm methodology (Katoch, Chauhan and Kumar, 2021; Kramer and Kramer, 2017).

The use of evolutionary optimization is particularly suitable when multiple objectives conflict. Increasing ensemble size may improve discrimination but simultaneously increase inference cost. Similarly, highly complex models may provide stronger classification boundaries but produce more difficult explanations. The optimization problem should therefore be formulated as a constrained multi-objective problem rather than as a simple accuracy maximization task.

3.4 Multi-Class Intrusion Decision

The framework is designed for multi-class detection rather than binary benign/malicious classification. A multi-class formulation is more useful in cyber-physical environments because different attack categories can produce different operational consequences. The ensemble generates a probability distribution over candidate classes, and the final decision is determined using adaptive weighted aggregation.

The architecture is conceptually aligned with the multi-class explainable ensemble direction demonstrated by Islam, Ahmed, Ishtiaq et al. (2026). However, adaptive weighting provides an additional mechanism for changing the influence of individual models when operating conditions vary.

3.5 Explainable AI Layer

Explainability is incorporated after the ensemble decision rather than being treated as an unrelated visualization feature. For every security decision, the system should identify the principal input characteristics contributing to the prediction and indicate the relative influence of the contributing classifiers.

An explanation can therefore be represented conceptually as:

E(x)={(f1,a1),(f2,a2),…,(fk,ak)}E(x)=\{(f_1,a_1),(f_2,a_2),\ldots,(f_k,a_k)\}

where fif_i represents an influential feature and aia_i its contribution to the prediction.

For example, if a drone communication instance is classified as malicious because several traffic characteristics substantially differ from its learned operational pattern, the explanation should expose those influential characteristics. In an industrial environment, an operator can then distinguish a security-relevant deviation from an ordinary operational change.

Explainability also has an important validation function. If the ensemble repeatedly identifies irrelevant variables as decisive, the explanation layer can expose potential model instability. Consequently, XAI is treated not merely as a user-interface mechanism but as a model-auditing component.

3.6 Operational Adaptation

The adaptive layer periodically evaluates recent detection behavior. If classifier performance changes, ensemble weights can be recalculated. A drift-sensitive strategy may therefore be represented as:

wm(t+1)=Φ(wm(t),Rm(t),Dt)w_m(t+1)=\Phi(w_m(t),R_m(t),D_t)

where Rm(t)R_m(t) represents recent reliability of classifier mm, DtD_t represents the current operational context, and Φ\Phi denotes the adaptation mechanism.

This design is especially relevant to heterogeneous systems because operational conditions are not necessarily stationary. A detector that remains unchanged after significant environmental changes may gradually lose effectiveness.

3.7 Computational Considerations

Adaptation introduces additional computational overhead. Genetic optimization requires population evaluation, ensemble inference requires multiple model executions, and explanation generation introduces another processing stage. Therefore, the architecture must balance detection quality against latency.

The high-throughput considerations discussed by Ferraz et al. (2021) and Shao et al. (2019) provide an important conceptual analogy: algorithmic performance cannot be evaluated independently of implementation cost. In a real-time cyber-physical environment, excessive latency may itself reduce the practical value of a detection mechanism.

A practical implementation should therefore use lightweight base learners where possible, cache model outputs when appropriate, restrict optimization to defined update intervals, and prioritize high-risk traffic for deeper analysis.

RESULTS

Because the supplied material does not provide a common experimental dataset, benchmark scores, or measured implementation results for the proposed architecture, the findings are presented as analytical and framework-level findings rather than fabricated empirical statistics. This distinction is essential for research integrity.

The first analytical finding is that adaptive ensemble detection provides a stronger theoretical basis for heterogeneous environments than a permanently fixed single-model detector. The reason is structural: different classifiers can specialize in different

behavioral patterns, while adaptive weighting permits their influence to change according to observed conditions. This principle is compatible with the broader literature on optimization and configuration search (Katoch, Chauhan and Kumar, 2021; Kramer and Kramer, 2017).

The second finding concerns evolutionary optimization. Genetic algorithms provide a mechanism for searching combinations of classifiers, features, and weights when the optimal configuration cannot be analytically determined. Prior applications of genetic optimization to coding-system design demonstrate that evolutionary search can identify high-performing configurations within technically constrained solution spaces (Elkelesh, Ebada, Cammerer and Ten Brink, 2019; Elkelesh, Ebada, Cammerer, Schmalen and Ten Brink, 2019). In the proposed framework, this capability can be redirected toward ensemble configuration.

The third finding is that explainability changes the operational value of intrusion detection. A classification decision without interpretable evidence may be difficult for an operator to validate, particularly when false positives can disrupt legitimate cyber-physical operations. The directly relevant work on explainable ensemble intrusion detection in heterogeneous drone and industrial networks supports the importance of explainability in this application domain (Islam, Ahmed, Ishtiaq et al., 2026).

The fourth finding is that computational complexity represents a fundamental trade-off. Increasing ensemble diversity and optimization frequency can potentially improve adaptability, but it also increases inference and training costs. High-throughput architecture studies in communication systems reinforce the importance of considering implementation efficiency alongside algorithmic performance (Ferraz et al., 2021; Shao et al., 2019).

Finally, the literature supports the feasibility of applying optimization concepts to industrial security. Zhou and Shi (2024) specifically connect genetic algorithms with industrial-control firewall whitelisting, indicating that evolutionary approaches can be relevant to adaptive security policy construction. Consequently, the proposed framework occupies a logically defensible position at the intersection of adaptive optimization, ensemble detection, explainability, and cyber-physical security.

DISCUSSION

The principal theoretical contribution of the proposed framework is the integration of three traditionally distinct concerns: classification diversity, adaptive optimization, and decision explainability. Ensemble learning addresses the weakness of relying on a single decision boundary, evolutionary optimization addresses the challenge of selecting an effective configuration, and XAI addresses the interpretability problem. Their integration is particularly meaningful for cyber-physical environments because security decisions may have consequences beyond conventional data loss.

The framework also highlights a fundamental contradiction between model sophistication and operational responsiveness. A larger ensemble may provide greater representational diversity, but each additional model increases inference requirements. Similarly, frequent genetic optimization may enable faster adaptation but consume computational resources. The high-throughput hardware literature demonstrates why this trade-off cannot be ignored when algorithms are deployed in resource-sensitive environments (Ferraz et al., 2021; Shao et al., 2019).

Another important issue concerns the reliability of adaptation. Adaptive weighting can respond to changing conditions, but adaptation based on unreliable recent observations could reinforce temporary anomalies or contaminated data. Consequently, adaptation

should not automatically treat every short-term performance change as evidence of genuine environmental drift. Robust validation, minimum update intervals, and safeguards against unstable weight oscillation would be necessary in a practical implementation.

Explainability introduces another trade-off. Highly complex ensemble decisions can be difficult to explain consistently. An explanation that identifies one dominant feature while the final decision is actually distributed across many classifiers may give an incomplete representation of the model's reasoning. Therefore, explanation quality should be evaluated not only by visual simplicity but also by fidelity to the actual ensemble decision. This issue is especially important in the context of explainable multi-class intrusion detection (Islam, Ahmed, Ishtiaq et al., 2026).

The provided literature also indicates that communication constraints cannot be separated from security architecture. Research on short-length error-correcting codes and communication-system design emphasizes performance under constrained conditions (Van Wonterghem et al., 2016; Rao, 2015). Drone networks may experience changing communication characteristics, making computational efficiency and communication robustness important components of a complete security design.

From an industrial perspective, the framework has potential value because adaptive security policies can complement intrusion detection. Zhou and Shi (2024) demonstrate the relevance of genetic algorithms to industrial-control firewall whitelisting. An adaptive IDS could therefore operate as part of a broader layered security mechanism in which detection identifies suspicious behavior and policy mechanisms determine appropriate containment.

Nevertheless, several limitations remain. The proposed methodology has not been validated here through a unified empirical benchmark, and the provided references do not establish a common dataset covering all intended drone and industrial attack categories. Consequently, numerical claims about accuracy, precision, recall, latency, or false-positive reduction would be unsupported. Future empirical evaluation should therefore compare static and adaptive ensembles under identical conditions, measure explanation fidelity, evaluate computational overhead, and test resilience under changing traffic distributions.

CONCLUSION

This paper proposed an adaptive ensemble-based intrusion detection framework for secure drone and industrial cyber-physical systems. The architecture integrates heterogeneous classifiers, adaptive weighting, genetic-algorithm-based configuration optimization, multi-class decision-making, and explainable artificial intelligence. Its theoretical foundation is supported by literature on genetic algorithms, optimization-oriented communication-system design, high-throughput architectures, industrial security optimization, and explainable ensemble intrusion detection.

The analysis also demonstrates that improved adaptability cannot be considered independently from computational cost and explanation reliability. A practical deployment must therefore optimize not only detection performance but also latency, resource consumption, stability, and interpretability. The research direction is particularly relevant to heterogeneous drone and industrial environments, where network behavior and operational requirements may differ substantially.

Future research should implement the framework on representative heterogeneous datasets, evaluate multi-class attack detection under distribution shifts, compare alternative adaptive weighting mechanisms, quantify explanation fidelity, and investigate

hardware-aware deployment strategies. Such empirical validation would allow the proposed architecture to progress from a theoretically grounded framework to a fully benchmarked cyber-physical intrusion detection system.

REFERENCES

1. Ebada, M., Elkelesh, A., Cammerer, S. and ten Brink, S. (2018, May) Scattered EXIT charts for finite length LDPC code design (In 2018 IEEE International Conference on Communications (ICC)), 1-7.

2. Elkelesh, A., Ebada, M., Cammerer, S., Schmalen, L. and Ten Brink, S. (2019) Decoder-in-the-loop: Genetic optimization-based LDPC code design (IEEE access), 7, 141161-141170.

3. Elkelesh, A., Ebada, M., Cammerer, S. and Ten Brink, S. (2019) Decoder-tailored polar code design using the genetic algorithm (IEEE Transactions on Communications), 67(7), 4521-4534.

4. Eremia, M., Liu, C. C. and Edris, A. A. (2016) Advanced solutions in power systems: HVDC, FACTS, and Artificial Intelligence (John Wiley & Sons).

5. Ferraz, O., Subramaniyan, S., Chinthala, R., Andrade, J., Cavallaro, J. R., Nandy, S. K. and Falcao, G. (2021) A survey on high-throughput non-binary LDPC decoders: ASIC, FPGA, and GPU architectures (IEEE Communications Surveys & Tutorials), 24(1), 524-556.

6. Katoch, S., Chauhan, S. S. and Kumar, V. (2021) A review on genetic algorithm: past, present, and future (Multimedia tools and applications), 80, 8091-8126.

7. Kramer, O. and Kramer, O. (2017) Genetic algorithms (Springer International Publishing).

8. Rao, K. D. (2015) Channel coding techniques for wireless communications (Berlin, Germany: Springer India).

9. Shao, S., Hailes, P., Wang, T. Y., Wu, J. Y., Maunder, R. G., Al-Hashimi, B. M. and Hanzo, L. (2019) Survey of turbo, LDPC, and polar decoder ASIC implementations (IEEE Communications Surveys & Tutorials), 21(3), 2309-2333.

10. Van Wonterghem, J., Alloum, A., Boutros, J. J. and Moeneclaey, M. (2016, November) Performance comparison of short-length error-correcting codes (In 2016 Symposium on Communications and Vehicular Technologies (SCVT)), 1-6.

11. Wijekoon, V. B., Viterbo, E., Hong, Y., Micheloni, R. and Marelli, A. (2019) A novel graph expansion and a decoding algorithm for NB-LDPC codes (IEEE Transactions on Communications), 68(3), 1358-1369.

12. Zhou, X., and Shi, W. (2024) Research on the optimisation of whitelisting technology for network firewall in industrial control system using genetic algorithm (International Journal of Communication Networks and Distributed Systems), 30(1), 30-41.

13. Islam, M.S., Ahmed, F., Ishtiaq, W. et al. Advanced explainable ensemble models for multi-class intrusion detection in heterogeneous drone and industrial networks. J. Inf. Secur. 2026, 14 (2026).

14. K. Ramamurthy, N. Bellamkonda and N. Amanmadov, "ScalePulse: Combinatorial Llm Framework for Scalability Constraint," SoutheastCon 2026, Huntsville, AL, USA, 2026, pp. 1-6, doi: 10.1109/SoutheastCon63549.2026.11476075

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