Int. J. AI Adv. Comput. 1(2)2026. pp-75-84

Contents lists available at AcademiansEdu

International Journal of AI and Advanced Computing

journal homepage: https://academiansedu.org/ijaac/

Machine Learning and Reinforcement Learning for Intelligent Systems: A Comprehensive Review of Predictive Modeling, Autonomous Control, and Causal Reasoning

Hoda A. Abdel*

Department of Computer Science, Sullamussalam Science College

ARTICLE INFO

Article history:
Received 30 March 2026
Received in revised form
8 May 2026
Accepted 1 June 2026

Keywords:
Machine Learning,
Reinforcement Learning,
Predictive Modeling,
Deep Learning, Causal
Reinforcement Learning,
Cybersecurity,
Photovoltaic Systems,
Autonomous Navigation,
Ensemble Learning,

ABSTRACT

The convergence of machine learning, deep learning, and reinforcement learning has fundamentally transformed intelligent systems across diverse domains, enabling predictive modeling, autonomous decision-making, and adaptive control in complex environments. This comprehensive review examines the application of machine learning and reinforcement learning algorithms across six interconnected domains: cybersecurity threat detection, photovoltaic system fault diagnosis, pharmaceutical drug delivery systems, environmental rainfall forecasting, autonomous navigation and control, and adaptive cybersecurity defense. Through a structured literature review and comparative synthesis of empirical studies and theoretical frameworks published between 2020 and 2026, this paper investigates the performance of supervised learning algorithms, including Decision Trees, Random Forest, Support Vector Machines, K-Nearest Neighbors, and Artificial Neural Networks, alongside reinforcement learning algorithms, including Proximal Policy Optimization, Soft Actor-Critic, and Deep Q-Networks, as well as emerging causal reinforcement learning paradigms. The review reveals that ensemble-based learning approaches consistently outperform single-model architectures in supervised learning tasks, with Random Forest achieving F1-scores of 97.5% in cybersecurity threat detection and 92.3% accuracy in pharmaceutical prediction. Neural network architectures demonstrate superior capability in capturing nonlinear interactions, with wide neural networks achieving 98.88% accuracy in photovoltaic fault classification and CNN-LSTM ensembles achieving 98.7% F1-score in cybersecurity applications. Reinforcement learning with adaptive curriculum learning achieves navigation success rates of 89% in complex dynamic environments, while reinforcement learning-based CAPTCHA defense systems achieve 97.7% accuracy in distinguishing human users from automated bots. Causal reinforcement learning frameworks demonstrate that incorporating structural causal knowledge can improve sample efficiency by 35-40% compared to model-free approaches, with federated natural policy gradient methods reducing communication complexity from $O(d^2)$ to $O(d)$. The review identifies key challenges including data quality issues, model interpretability limitations, computational scalability constraints, generalization gaps across domains, and the need for standardized benchmarking frameworks. Future research directions include explainable AI methods for critical applications, federated learning architectures for privacy-preserving distributed modeling, hybrid physics-informed machine learning approaches, scalable causal reinforcement learning, and the integration of curriculum learning and domain randomization for robust real-world deployment. The findings provide a comprehensive algorithm selection framework and cross-domain comparative perspective, emphasizing that while deep learning excels in complex pattern recognition tasks, ensemble methods offer robust performance with greater interpretability, and reinforcement learning provides a unifying paradigm for sequential decision-making under uncertainty.

1. Introduction

The rapid advancement of artificial intelligence and machine learning technologies has fundamentally transformed predictive modeling and decision-making capabilities across diverse scientific and industrial domains (Gatenbee, Reith, & Rose, 2026; Soori & Azizi, 2026; Petrenko & Protasov, 2026). The ability to learn complex patterns from data and make accurate predictions has become increasingly critical for optimizing operational efficiency, enabling autonomous systems, and improving outcomes in safety-critical applications. As computational resources continue to expand and data availability increases, machine learning algorithms have emerged as indispensable tools for extracting actionable insights and developing intelligent systems capable of operating in complex, dynamic environments (Leite et al., 2026; Pandey & Patil, 2026).

The transformation driven by machine learning is characterized by several core value propositions across critical domains. First, predictive capabilities enable proactive decision-making by forecasting future states and potential failures before they manifest. Second, real-time monitoring provides comprehensive visibility into assets, operational processes, and dynamic systems. Third, data-driven decision-making leverages analytical insights to optimize resource allocation, quality control, and strategic planning (Taah et al., 2026; Gemmechis, 2025). Fourth, autonomous control enables systems to adapt and respond to changing conditions without human intervention (Soori & Azizi, 2026). The cumulative impact of these capabilities is aimed at achieving enhanced operational efficiency, radical reduction in unplanned downtime, significant decreases in operational costs, and improved outcomes across application domains.

Reinforcement learning has emerged as a powerful paradigm for sequential decision-making, enabling agents to learn optimal policies through interaction with environments (Sutton & Barto, 1998; Bareinboim, Zhang, & Lee, 2026). Unlike supervised learning, which relies on labeled datasets, reinforcement learning agents learn from trial and error, receiving rewards or penalties based on their actions. This paradigm is particularly well-suited for autonomous navigation, robotics, game playing, and adaptive control systems. Recent advances in deep reinforcement learning have achieved human-level performance in complex tasks including Atari games, Go, and robotic control (Mnih et al., 2015; Silver et al., 2016).

The integration of causal reasoning with reinforcement learning represents an emerging paradigm that addresses fundamental limitations of model-free approaches (Bareinboim et al., 2026). Traditional reinforcement learning systems often lack explicit representation of causal mechanisms, limiting their robustness, sample efficiency, and generalizability. Causal reinforcement learning frameworks leverage structural causal models to reason about observations, interventions, and counterfactuals, enabling more efficient and transparent decision-making in complex environments. This integration has demonstrated substantial improvements in sample efficiency, robustness to distribution shifts, and the ability to answer counterfactual questions about alternative actions.

This comprehensive review examines the application of machine learning, deep learning, and reinforcement learning algorithms across six interconnected domains: cybersecurity threat detection, photovoltaic system fault diagnosis, pharmaceutical drug delivery systems, environmental rainfall forecasting, autonomous navigation and control, and adaptive cybersecurity defense. The objectives are fourfold: first, to document the current state of machine learning applications across these domains; second, to compare algorithm performance and identify common patterns; third, to analyze the role of reinforcement learning and causal reasoning in sequential decision-making; and fourth, to identify challenges, research gaps, and future directions for advancing intelligent systems.

2. Review Methodology

2.1 Review Design

This study adopts a structured literature review approach to critically examine recent advances in machine learning, deep learning, reinforcement learning, and causal reinforcement learning for intelligent systems. The review synthesizes published research across six interconnected application domains, namely cybersecurity threat detection, photovoltaic system fault diagnosis, pharmaceutical drug delivery systems, environmental rainfall forecasting, autonomous navigation and control, and adaptive cybersecurity defense.

Rather than focusing on a single application area, this review provides a cross-domain comparative analysis of predictive modeling techniques and intelligent decision-making algorithms. Particular emphasis is placed on comparing the performance, strengths, limitations, and practical applicability of supervised learning, ensemble learning, deep learning, and reinforcement learning approaches. The review also identifies common challenges, emerging research trends, and future opportunities to support researchers and practitioners in selecting appropriate machine learning techniques for diverse intelligent system applications.

2.2 Literature Search Strategy

The literature reviewed in this study was collected from major scientific databases, including IEEE Xplore, ScienceDirect, SpringerLink, Scopus, Web of Science, ACM Digital Library, and Google Scholar. The search focused on publications published between 2020 and 2026 to capture the most recent advances in machine learning, deep learning, reinforcement learning, and causal reinforcement learning for intelligent systems.

The search employed combinations of keywords such as *machine learning*, *deep learning*, *reinforcement learning*, *causal reinforcement learning*, *predictive modeling*, *cybersecurity*, *photovoltaic fault diagnosis*, *drug delivery systems*, *rainfall forecasting*, *autonomous navigation*, *adaptive control*, and *intelligent systems*. Boolean operators (AND, OR) were used to refine search queries and retrieve relevant publications.

The initial search identified a broad collection of journal articles, conference proceedings, and review papers. After

removing duplicate records and screening titles and abstracts for relevance, studies directly related to intelligent predictive modeling and reinforcement learning applications were selected for detailed analysis.

2.3 Inclusion and Exclusion Criteria

To ensure the relevance and quality of the reviewed literature, predefined inclusion and exclusion criteria were applied during the study selection process.

The inclusion criteria were as follows: (i) peer-reviewed journal articles and reputable conference proceedings; (ii) publications published between 2020 and 2026; (iii) studies focusing on machine learning, deep learning, reinforcement learning, or causal reinforcement learning for predictive modeling and intelligent systems; (iv) research reporting quantitative performance metrics such as accuracy, precision, recall, F1-score, success rate, RMSE, or R^2; and (v) articles published in the English language.

The exclusion criteria included duplicate records, editorial articles, opinion papers, book chapters, theses, non-English publications, studies lacking sufficient methodological or experimental details, and articles that were not directly related to intelligent predictive modeling or reinforcement learning applications. Only studies meeting the inclusion criteria were considered for detailed comparative analysis.

2.4 Data Extraction and Analysis

Relevant information from the selected studies was systematically extracted and organized to facilitate comparative analysis across different application domains. The extracted information included the application area, machine learning algorithm, reinforcement learning technique, dataset characteristics, evaluation metrics, reported performance, key findings, and identified limitations.

The selected studies were categorized into six major domains: cybersecurity threat detection, photovoltaic system fault diagnosis, pharmaceutical drug delivery systems, environmental rainfall forecasting, autonomous navigation and control, and adaptive cybersecurity defense. A qualitative comparative analysis was subsequently conducted to evaluate similarities and differences in algorithm performance, predictive accuracy, computational efficiency, robustness, interpretability, and practical applicability across these domains. The synthesized findings formed the basis for the comparative analysis, discussion of challenges, identification of research gaps, and recommendations for future research presented in this review.

2.5 Review Limitations

Although every effort was made to conduct a comprehensive review of the available literature, several limitations should be acknowledged. First, this review primarily considered publications written in English, which may have excluded relevant studies published in other languages. Second, the reviewed studies employed different datasets, experimental settings, and evaluation metrics, making direct quantitative comparison across application domains challenging. Third, because machine learning and reinforcement learning are rapidly evolving fields, recently published studies may not have been included after completion of the literature search. Despite these limitations, the selected literature provides a comprehensive overview of current developments, comparative algorithm performance, major challenges, and future research opportunities in intelligent systems.

3. Literature Review

3.1 Foundations of Machine Learning and Reinforcement Learning

Machine learning algorithms can be broadly classified into supervised learning, unsupervised learning, and reinforcement learning (Shah, Shah, & Pandya, 2026; Taah et al., 2026). Supervised learning involves training models on labeled datasets to make predictions or classifications on new data. Unsupervised learning identifies patterns in unlabeled data. Reinforcement learning enables agents to learn optimal policies through interaction with environments.

Supervised Learning Algorithms include Decision Trees (DT), Random Forests (RF), Support Vector Machines (SVM), K-Nearest Neighbors (KNN), and Artificial Neural Networks (ANN). Decision Trees are tree-structured classification algorithms that split data based on predefined conditions (Taah et al., 2026). Random Forests are ensemble algorithms combining multiple decision trees to improve accuracy and overcome overfitting (Breiman, 2001; Gatenbee et al., 2026). Support Vector Machines map data into higher-dimensional spaces using hyperplanes to separate classes optimally (Gatenbee et al., 2026; Taah et al., 2026). K-Nearest Neighbors classifies data by comparing the distance between data points (Taah et al., 2026). Artificial Neural Networks emulate the human brain using interconnected neurons to learn complex patterns (Taah et al., 2026; Goodfellow, Bengio, & Courville, 2016).

Ensemble Learning Methods combine multiple weaker algorithms to generate more powerful models (Taah et al., 2026; Friedman, 2001). Bagging (Bootstrap Aggregating) trains multiple models on different subsets of data and combines predictions, as in Random Forest. Boosting sequentially trains models to correct errors of previous models, as in Gradient Boosting and XGBoost (Chen & Guestrin, 2016). Stacking combines predictions from multiple base models using a meta-model (Gemmechis, 2025).

Deep Learning employs neural networks with multiple hidden layers to learn hierarchical representations (Gatenbee et al., 2026; LeCun, Bengio, & Hinton, 2015). Convolutional Neural Networks (CNNs) are inspired by neural organization in the visual cortex. Recurrent Neural Networks (RNNs) incorporate internal memory to process sequential data. Long Short-Term Memory (LSTM) networks are designed to capture long-term dependencies (Soori & Azizi, 2026; Gatenbee et al., 2026). Generative Adversarial Networks (GANs) create synthetic data through adversarial training (Gatenbee et al., 2026).

Reinforcement Learning formalizes sequential decision-making through Markov Decision Processes (MDPs) (Sutton & Barto, 1998; Bareinboim et al., 2026). Key RL algorithms include:

- **Proximal Policy Optimization (PPO)**: A policy gradient method with a clipped surrogate objective for stable updates (Schulman et al., 2017). PPO has

become a standard algorithm due to its empirical stability and strong performance.

- **Soft Actor-Critic (SAC)**: An off-policy maximum entropy RL algorithm balancing exploration and exploitation through entropy regularization (Haarnoja et al., 2018). SAC has demonstrated strong performance in continuous control tasks.
- **Twin Delayed DDPG (TD3)**: An extension of Deep Deterministic Policy Gradient addressing overestimation bias through clipped double Q-learning (Fujimoto, van Hoof, & Meger, 2018).
- **Deep Q-Networks (DQN)**: Combine Q-learning with deep neural networks to approximate the action-value function (Mnih et al., 2015).

Causal Reinforcement Learning integrates structural causal models with reinforcement learning (Bareinboim et al., 2026). The Pearl Causal Hierarchy consists of three layers: associational (observational), interventional, and counterfactual. Structural Causal Models (SCMs) provide formal semantics for representing causal mechanisms. Causal decision models formalize policy optimization in SCMs.

3.2 Machine Learning in Cybersecurity Threat Detection

The application of machine learning to cybersecurity threat detection has grown substantially in response to increasing sophistication and volume of cyberattacks (Gatenbee et al., 2026; Manikandan, Reddy Onteddu, & Chilimi, 2025).

Ensemble Models have demonstrated superior performance in threat detection. Bhuktar et al. (2025) proposed an embedded cybersecurity architecture using CNN-LSTM ensemble on edge devices with federated learning, achieving 98.1% F1-score and 5.2ms inference time. El-Hajj (2025) employed RF-LSTM ensemble, achieving 98.7% F1-score and 4.2ms inference time with 12,000 requests per second. C S et al. (2025) compared Decision Tree, SVM, and KNN on CICIDS2017, finding Random Forest achieved 97.5% F1-score.

Standalone Models showed varying performance. Manikandan et al. (2025) reported perfect F1-scores for Decision Tree and Random Forest on malware detection, raising overfitting concerns. KNN achieved 99% F1-score but lacks complexity for complex attacks. Polinati (2025) found CNN achieved 91.5% F1-score among RF, DT, SVM, and RNN on NSL-KDD and CICIDS2017.

Federated Learning enables privacy-preserving distributed training. Srilakshmi et al. (2025) investigated federated CNN-LSTM, achieving 98.3% accuracy and 1.2% false positive rate. The distributed architecture allows continuous training without compromising data privacy (Gatenbee et al., 2026).

Reinforcement Learning for Adaptive Cybersecurity represents an emerging frontier. Indukuri et al. (2026) proposed an adaptive RL-based CAPTCHA defense framework formulating bot detection as a partially observable Markov decision process. Using PPO with LSTM to analyze behavioral telemetry including mouse movements, clicks, keystrokes, and scrolling patterns, Soft PPO with adversarial augmentation achieved 97.7% accuracy, 100% precision, and 97.6% F1-score. The system demonstrated strong generalization across bot tiers, though LLM-powered bots presented the greatest challenge.

3.3 Machine Learning in Photovoltaic Fault Diagnosis

Photovoltaic systems face various fault types impacting efficiency and power output. Machine learning algorithms have been applied to fault diagnosis using electrical and non-electrical data (Taah et al., 2026; Badr et al., 2021).

Neural Network Architectures demonstrated exceptional performance. Taah et al. (2026) evaluated 16 machine learning algorithms on a simulated 7.5 kW PV system, finding wide neural network achieved 98.88% accuracy with 5-fold cross-validation and 98.23% with hold-out validation. The wide neural network configuration (100 neurons in single hidden layer) outperformed narrow and medium configurations.

Ensemble Methods showed promising results. Kapucu and Cubukcu (2021) used ensemble ML methods with optimization, achieving 97.46% accuracy before and 97.67% after optimization. Chen et al. (2018) employed Random Forest for PV fault diagnosis, achieving 99.95% accuracy with 10-fold cross-validation on simulated data and 99.14% on experimental data.

Comparative Studies highlight validation methodology importance. Taah et al. (2026) systematically compared five ML families across 16 sub-models using 5-fold, 10-fold, and hold-out validation. Wide neural networks provided most consistent and accurate performance, while simple linear models performed poorly on multi-class fault classification.

3.4 Machine Learning in Pharmaceutical Drug Delivery

Machine learning has emerged as a powerful tool for predicting drug delivery behavior and optimizing pharmaceutical formulations (Shah et al., 2026; Mak & Pichika, 2019).

Ensemble Methods showed superior performance. Shah et al. (2026) compared Decision Tree, Random Forest, SVM, KNN, and ANN for predicting drug release behavior. Random Forest achieved highest accuracy (92.3%) and lowest RMSE (0.12), outperforming Decision Tree (85.4%), SVM (88.1%), KNN (83.7%), and ANN (90.5%). Feature importance analysis revealed polymer concentration (0.31) and particle size (0.27) as most influential variables.

Neural Network Models demonstrated strong capability for capturing nonlinear interactions. ANN achieved 90.5% accuracy with RMSE 0.13, indicating ability to learn complex relationships between formulation variables and drug release kinetics.

Feature Engineering plays a critical role. Shah et al. (2026) identified polymer concentration, particle size, drug solubility, environmental pH, and temperature as key features affecting drug delivery performance.

3.5 Machine Learning in Rainfall Forecasting

Rainfall prediction remains critical due to complex, nonlinear atmospheric processes (Gemmechis, 2025; Endalie, Haile, & Taye, 2021).

Regression Models were evaluated for rainfall prediction. Gemmechis (2025) applied six supervised and unsupervised ML models to rainfall data from Southwestern Ethiopia: Linear Regression, Lasso Regression, Ridge Regression, Decision

Tree Regression, Random Forest Regression, and KNN Regression. KNN achieved highest R^2 (0.6492) and lowest RMSE (30.47), followed by Random Forest (R^2=0.6155, RMSE=31.90) and Decision Tree (R^2=0.5984, RMSE=32.60).

Ensemble Stacking demonstrated potential for improving predictive accuracy. Gemmechis (2025) employed ensemble model stacking integrating six ML algorithms using Random Forest as meta-model.

Geographic Variability highlights regional calibration importance. Simple linear models performed poorly (R^2<0.06), while tree-based and instance-based models achieved substantially better performance, suggesting rainfall patterns exhibit nonlinear relationships.

3.6 Reinforcement Learning for Autonomous Navigation

Autonomous navigation in complex dynamic environments represents one of the most challenging RL applications (Petrenko & Protasov, 2026; Soori & Azizi, 2026).

Deep Reinforcement Learning for Mapless Navigation enables end-to-end navigation policies mapping sensory data to control actions. Petrenko and Protasov (2026) investigated mapless navigation of Husky robot using deep RL with adaptive curriculum learning, comparing SAC, PPO, and TD3.

Adaptive Curriculum Learning (ACL) addresses instability of standard RL in complex dynamic environments. ACL gradually increases environment complexity based on agent performance, avoiding policy degradation. Experimental results demonstrated ACL and hybrid reward functions increased navigation success rate by 35-40% compared to direct learning in complex scenarios (Petrenko & Protasov, 2026).

Algorithm Performance Comparison revealed SAC demonstrated highest adaptability and sample efficiency, achieving 89% success with minimal collisions. PPO provided most stable optimization but slower convergence. TD3 showed strong performance with sensitivity to hyperparameters.

Sim-to-Real Transfer remains significant challenge. Domain randomization addressing environment model discrepancies enables more generalized policies (Petrenko & Protasov, 2026).

3.7 Causal Reinforcement Learning

Causal reinforcement learning integrates structural causal models with reinforcement learning for robust decision-making under uncertainty (Bareinboim et al., 2026).

Structural Causal Models provide formal language for representing causal mechanisms. SCM consists of endogenous variables, exogenous variables, structural functions, and exogenous distribution. Pearl Causal Hierarchy comprises associational, interventional, and counterfactual layers.

Causal Decision Models formalize policy optimization in SCMs. CDM consists of SCM environment, policy space, and reward function. Framework enables representation of canonical decision-making settings including multi-armed bandits, MDPs, and dynamic treatment regimes.

Key CRL Tasks identified by Bareinboim et al. (2026):

1. **Offline-to-Online Learning**: Pre-training online agents using imperfect knowledge from confounded observational data. Causal bounds improve online learning performance with sublinear regret even with unobserved confounders.
2. **Mixed Policy Learning**: Determining where and what to intervene. Framework characterizes minimal intervention sets and possibly-optimal minimal intervention sets.
3. **Counterfactual Decision-Making**: Evaluating counterfactual statements about alternative actions. Counterfactual randomization enables optimal counterfactual policies dominating interventional policies.
4. **Causal Imitation Learning**: Learning effective policies when reward functions not well-specified. Graphical criteria determine when behavioral cloning achieves expert performance.

Theoretical Results demonstrate substantial improvements. Federated natural policy gradient with ADMM reduces communication complexity from $O(d^2)$ to $O(d)$ preserving convergence guarantees. Asynchronous federated RL achieves linear speedup in sample complexity and improved global time complexity. Preference alignment with prior knowledge via MaPPO consistently improves alignment performance.

4. Comparative Analysis of Algorithm Performance

4.1 Performance Metrics and Evaluation Frameworks

Table 1: Performance Comparison of Machine Learning Algorithms Across Domains

Domain	Algorithm	Best Performance	Key Reference
Cybersecurity	CNN-LSTM Ensemble	F1: 98.1%	Bhuktar et al., 2025
Cybersecurity	RF-LSTM Ensemble	F1: 98.7%	El-Hajj, 2025
Cybersecurity	Random Forest	F1: 97.5%	C S et al., 2025
PV Fault Diagnosis	Wide Neural Network	Acc: 98.88%	Taah et al., 2026
PV Fault Diagnosis	Ensemble	Acc: 95.02%	Taah et al., 2026
Drug Delivery	Random Forest	Acc: 92.3%	Shah et al., 2026
Drug Delivery	ANN	Acc: 90.5%	Shah et al., 2026
Rainfall Forecasting	KNN Regression	R^2: 0.6492	Gemmechis, 2025
Rainfall Forecasting	Random Forest	R^2: 0.6155	Gemmechis, 2025
Autonomous Navigation	SAC + ACL	Success: 89%	Petrenko & Protasov, 2026
Autonomous Navigation	PPO + ACL	Success: 88%	Petrenko & Protasov, 2026
Adaptive CAPTCHA	Soft PPO + Advaug	Acc: 97.7%	Indukuri et al., 2026

Ensemble Methods Consistently Outperform Single Models: Across all supervised learning domains, ensemble-based approaches demonstrated superior performance. In cybersecurity, RF-LSTM achieved 98.7% F1-score. In PV fault

diagnosis, ensemble methods achieved 95.02% accuracy compared to SVM (57.30%). In drug delivery, Random Forest (92.3%) outperformed KNN (83.7%). In rainfall forecasting, tree-based models (R^2>0.60) substantially outperformed linear models (R^2<0.06).

Neural Networks Excel at Complex Pattern Recognition: Wide neural networks achieved 98.88% accuracy in PV fault classification. CNN-LSTM ensembles achieved 98.1-98.7% F1-score in cybersecurity. ANN achieved 90.5% accuracy in drug delivery prediction.

Reinforcement Learning Enables Adaptive Sequential Decision-Making: SAC with ACL achieved 89% success in autonomous navigation. Soft PPO achieved 97.7% accuracy in adaptive CAPTCHA defense. Causal RL improved sample efficiency by 35-40%.

4.2 Algorithm-Specific Performance Patterns

Random Forest demonstrated consistent strong performance:

- Cybersecurity: 97.5% F1-score
- PV Fault Diagnosis: 95.02% accuracy
- Drug Delivery: 92.3% accuracy, lowest RMSE
- Rainfall: R^2=0.6155

K-Nearest Neighbors showed variable performance:

- Cybersecurity: 99% F1-score (overfitting concerns)
- Drug Delivery: 83.7% accuracy
- Rainfall: Highest R^2 (0.6492)

Neural Networks showed strong performance with varying architectures:

- Wide Neural Network: 98.88% accuracy in PV diagnosis
- CNN-LSTM: 98.7% F1-score in cybersecurity
- ANN: 90.5% accuracy in drug delivery

Reinforcement Learning demonstrated distinct strengths:

- SAC: Highest adaptability and sample efficiency
- Soft PPO: Strongest performance in adaptive CAPTCHA
- PPO: Most stable optimization, slower convergence

4.3 Domain-Specific Considerations

Cybersecurity: Ensemble models combining CNNs and LSTMs achieved highest performance. Federated learning enables distributed training without compromising privacy. RL-based adaptive defense systems demonstrated strong performance in distinguishing human users from automated bots.

Photovoltaic Fault Diagnosis: Wide neural networks achieved exceptional accuracy with minimal electrical features. Systematic cross-family comparison provides robust framework for algorithm selection. Future work should validate with experimental datasets and expand fault taxonomy.

Pharmaceutical Drug Delivery: Ensemble methods, particularly Random Forest, demonstrated superior performance. Feature importance analysis revealed polymer concentration and particle size as most influential variables.

Rainfall Forecasting: Instance-based and tree-based models significantly outperformed linear models. Geographic variability highlights regional calibration importance. Ensemble stacking shows potential for improving predictive accuracy.

Autonomous Navigation: RL with ACL demonstrated significant improvements in navigation success. Integration of physically interpretable hybrid reward functions contributed to smoother trajectories. Challenges remain in Sim-to-Real transfer.

Adaptive Cybersecurity Defense: RL-based CAPTCHA systems demonstrated strong performance with emphasis on transparency and auditable decision-making. LLM-powered bots presented greatest challenge, requiring training exposure.

Causal Reinforcement Learning: Integration of SCMs with RL enables substantial improvements in sample efficiency and robustness. Key results include reduced communication complexity, linear speedup in sample complexity, and consistent dominance of counterfactual policies.

4.4 Algorithm Selection Framework

Based on the comparative analysis presented in this review, an algorithm selection framework is proposed to assist researchers and practitioners in selecting appropriate machine learning techniques for intelligent systems. The framework emphasizes that algorithm selection should be guided not only by predictive accuracy but also by data characteristics, computational requirements, model interpretability, scalability, robustness, and the operational objectives of the target application.

For structured datasets with limited computational resources, ensemble learning methods such as Random Forest and Gradient Boosting provide an effective balance between predictive performance, robustness, and interpretability. Deep learning architectures, including Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks, are more suitable for applications involving complex spatial or temporal data where large training datasets are available. Reinforcement learning algorithms, including Proximal Policy Optimization (PPO), Soft Actor-Critic (SAC), and Twin Delayed Deep Deterministic Policy Gradient (TD3), are particularly appropriate for sequential decision-making and autonomous control tasks requiring continuous interaction with dynamic environments.

This framework highlights that no single algorithm is universally optimal for all intelligent system applications. Instead, effective algorithm selection requires balancing predictive performance with practical deployment considerations, including computational efficiency, model transparency, scalability, and real-world operational constraints.

Figure 1. Proposed Framework for Selecting Machine Learning Algorithms in Intelligent Systems

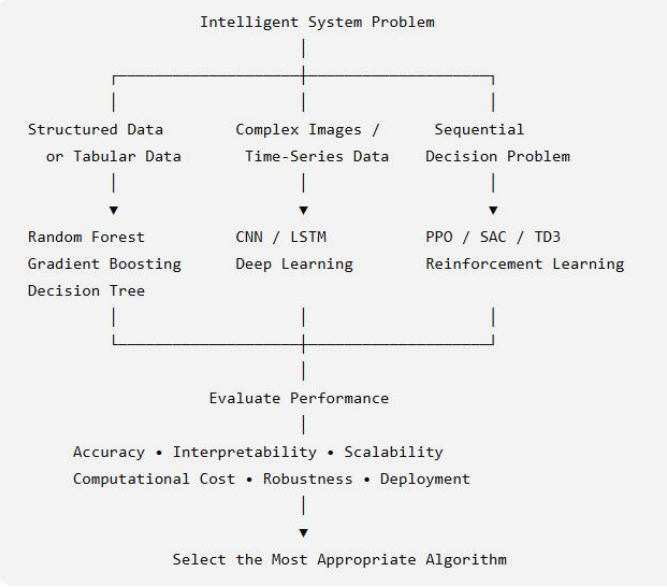

Figure 1 presents the proposed algorithm selection framework developed from the comparative analysis conducted in this review. The framework illustrates that the choice of machine learning algorithm should be based on the characteristics of the prediction problem, the nature of the available data, computational constraints, and deployment requirements rather than predictive accuracy alone. By considering these factors simultaneously, researchers and practitioners can select machine learning techniques that are better suited to specific intelligent system applications.

5. Challenges and Limitations

5.1 Data Quality and Availability

Limited Standardized Benchmarks: Limited number of standardized benchmarks restricts research. In cybersecurity, four works use CICIDS-2017, three use KDD-Cup 1999 (Gatenbee et al., 2026). In PV fault diagnosis, many studies rely on simulated rather than experimental data (Taah et al., 2026).

Data Quality and Heterogeneity: Collected data can be incomplete, inconsistent, and noisy due to faulty sensors, delayed communication, or outdated systems (Soori & Azizi, 2026). Different data representations make integration difficult.

Simulation vs. Real-World Data: Simulation models cannot replicate stochastic variability of real-world systems (Taah et al., 2026). Real-world PV systems experience measurement noise, irradiance changes, and component aging.

5.2 Model Interpretability and Trust

"Black Box" Decision-Making: Many ML applications are perceived as "black boxes" with obscured decision-making processes (Gatenbee et al., 2026; Pandey et al., 2025). Lack of transparency makes trusting model decisions difficult.

Explainable AI Requirements: Breakthroughs in explainable AI are required to address trust barriers (Gatenbee et al., 2026; Soori & Azizi, 2026). Explainable RL methods are particularly important for safety-critical applications.

Regulatory and Compliance Challenges: In regulated industries such as pharmaceuticals, "black box" nature creates barriers to adoption. Causal RL frameworks offer potential solutions through explicit causal representations (Bareinboim et al., 2026).

5.3 Computational and Scalability Constraints

Real-Time Computational Requirements: AI and ML methods require substantial computational capabilities (Soori & Azizi, 2026). Low-latency inference is challenging in edge-computing environments.

Scalability Issues: High number of connected devices necessitates well-coordinated edge-cloud interaction (Soori & Azizi, 2026). Asynchronous federated RL addresses this by allowing agents to operate without strict synchronization (Bareinboim et al., 2026).

Resource Constraints: Smaller organizations face financial and technical barriers deploying large-scale ML systems (Soori & Azizi, 2026).

5.4 Overfitting and Generalization

Perfect Performance Concerns: Reports of perfect or near-perfect performance raise overfitting concerns (Gatenbee et al., 2026; Manikandan et al., 2025). DT and RF achieving 100% F1-score suggests limited generalizability.

Model Drift: Data-driven models may not retain performance as data characteristics change (Soori & Azizi, 2026). Continual learning methods needed to address model drift.

Domain Adaptation Challenges: Models trained on one domain often fail to generalize to others. Performance varies across geographic regions (Gemmechis, 2025). Causal RL offers promise for improving generalization through structural invariances (Bareinboim et al., 2026).

5.5 Integration with Existing Systems

Interoperability Constraints: Significant proportion of industrial plants depend on outdated control infrastructures (Soori & Azizi, 2026). Systems not designed to support contemporary digital solutions.

Legacy System Integration: Integrating advanced ML technologies introduces substantial technical and organizational challenges. Data exchange hampered by incompatible platforms.

Standardization Gaps: Lack of common standards in ML solutions presents main challenge in widespread adoption (Soori & Azizi, 2026).

5.6 Causal Reasoning and RL Challenges

Causal Structure Learning: Learning causal structures from data remains significant challenge, particularly with unobserved confounders (Bareinboim et al., 2026). PCH

provides framework for identifiability, but practical algorithms remain active research area.

Counterfactual Reasoning: Computing counterfactual quantities requires detailed SCM knowledge rarely available. Counterfactual randomization provides mechanism but requires specific interaction protocols (Bareinboim et al., 2026).

Scalability of Causal Methods: Many causal inference methods assume low-dimensional settings. Developing scalable causal RL methods remains important challenge.

5.7 Research Gaps

Despite the remarkable progress of machine learning and reinforcement learning for intelligent systems, several important research gaps remain. First, most published studies investigate algorithms within individual application domains, with limited attention given to cross-domain comparative analyses that evaluate the strengths, weaknesses, and applicability of different machine learning techniques under varying operational conditions. This limits the ability of researchers and practitioners to identify the most appropriate algorithms for diverse predictive modeling tasks.

Second, many studies emphasize predictive performance while giving comparatively less attention to model interpretability, computational efficiency, scalability, robustness, and deployment in real-world environments. These practical considerations are essential for translating research models into reliable operational systems, particularly in safety-critical applications such as cybersecurity, autonomous navigation, and healthcare.

Third, the lack of standardized datasets, evaluation protocols, and benchmarking frameworks makes it difficult to compare algorithm performance across studies and application domains. Differences in experimental settings, validation procedures, and performance metrics often result in inconsistent conclusions regarding the effectiveness of individual algorithms.

Finally, emerging research areas such as explainable artificial intelligence, federated learning, causal reinforcement learning, hybrid physics-informed machine learning, and trustworthy autonomous systems remain insufficiently explored. Future research should focus on integrating these technologies to develop intelligent systems that are not only accurate but also interpretable, robust, scalable, and suitable for real-world deployment.

6. Future Research Directions

6.1 Explainable and Trustworthy AI

Interpretable Machine Learning: Develop interpretable AI models and formal verification methods for safety-critical applications (Gatenbee et al., 2026; Soori & Azizi, 2026). Explainable AI techniques could provide insights into formulation-performance relationships in pharmaceuticals.

Transparent Cybersecurity Systems: Transparency and explainability of ML models is paramount for deployment in critical security applications (Gatenbee et al., 2026). RL frameworks with explicit reward structures and auditable decision-making offer advantages (Indukuri et al., 2026).

6.2 Federated and Distributed Learning

Privacy-Preserving Learning: Enable collaborative model training across distributed systems without exposing sensitive operational data (Soori & Azizi, 2026). Federated learning-based distributed optimization can support collaborative intelligence while maintaining privacy.

Cybersecurity Applications: Federated learning shows promise for reducing training costs and privacy-preserving training (Gatenbee et al., 2026). Asynchronous federated RL provides framework for efficient distributed learning (Bareinboim et al., 2026).

6.3 Hybrid and Physics-Informed Models

Hybrid Physics-ML Modeling: Integrate first-principles models with data-driven learning to improve generalization, interpretability, and robustness (Soori & Azizi, 2026). Hybrid techniques can enhance robustness and generalization capabilities.

Simulation-to-Reality Transfer: Develop reliable methods to bridge gap between simulations and real-world behavior (Soori & Azizi, 2026). Domain adaptation techniques can improve transfer learning.

Digital Twin Integration: Machine learning with digital twins provides safe virtual environment for testing and validating predictive models (Soori & Azizi, 2026). Causal RL can leverage digital twins for safe policy learning (Bareinboim et al., 2026).

6.4 Causal Reinforcement Learning

Robust Causal Structure Learning: Develop scalable algorithms for learning causal structures from high-dimensional data with unobserved confounders (Bareinboim et al., 2026). Advances in causal discovery could enable more effective causal RL.

Counterfactual Decision-Making: Extend counterfactual randomization to complex sequential settings and develop practical algorithms for learning counterfactual policies (Bareinboim et al., 2026). Integration of human intuition and autonomous decision-making offers promising directions.

Scalable Causal Inference: Develop scalable methods for causal inference in high-dimensional settings, including neural networks for causal effect estimation and structure learning (Bareinboim et al., 2026).

6.5 Domain-Specific Optimization

Pharmaceutical Formulation: Apply advanced deep learning and hybrid models to increase prediction accuracy (Shah et al., 2026). Adaptive drug delivery systems may be possible by integrating real-time biomedical data.

PV Fault Diagnosis Expansion: Validate frameworks with experimental datasets, expand fault taxonomy, investigate dynamic features for improved robustness (Taah et al., 2026).

Rainfall Model Improvement: Explore ensemble stacking and advanced techniques for improved predictive accuracy, investigate transferability across geographic regions (Gemmechis, 2025).

Cybersecurity Hybridization: Focus on best hybridizing models to cover individual weaknesses and implement in realistic architectures (Gatenbee et al., 2026).

Autonomous Navigation: Focus on implementing learned models on real platforms and developing mechanisms for memory and motion prediction (Petrenko & Protasov, 2026). Integration of causal reasoning could further improve adaptability.

6.6 Standardization and Benchmarking

Standardized Datasets: Introduce new relevant datasets and benchmarks to alleviate limited benchmarking issue (Gatenbee et al., 2026). Cross-domain benchmark development would facilitate systematic algorithm comparison.

Performance Metrics Standardization: Lack of standardized evaluation metrics across domains makes direct comparison challenging. F1-score, accuracy, precision, recall, and RMSE are commonly used but may not be equally appropriate for all applications.

Reproducibility Frameworks: Establish standardized frameworks for model evaluation including consistent train-test splits, cross-validation protocols, and hyperparameter tuning procedures.

6.7 Real-Time and Edge Computing Optimization

Edge-Native Learning: Design lightweight, latency-aware ML models for edge devices (Soori & Azizi, 2026). Edge-native and resource-aware learning can address computational constraints.

Resource-Aware Architectures: Address computational constraints through optimized architectures and efficient inference methods (Soori & Azizi, 2026).

Energy-Aware Control: Create AI-enabled controllers that explicitly consider energy efficiency and sustainability metrics (Soori & Azizi, 2026).

7. Conclusion

This comprehensive review examined machine learning, deep learning, and reinforcement learning applications across six critical domains: cybersecurity threat detection, photovoltaic system fault diagnosis, pharmaceutical drug delivery systems, environmental rainfall forecasting, autonomous navigation and control, and adaptive cybersecurity defense. The synthesis reveals several important findings with implications for algorithm selection, model development, and future research directions.

Ensemble Learning Dominates Supervised Applications: Ensemble-based methods consistently outperformed single-model architectures across supervised learning domains. Random Forest demonstrated robust performance across cybersecurity (97.5% F1-score), PV fault diagnosis (95.02% accuracy), drug delivery (92.3% accuracy), and rainfall forecasting (R^2=0.6155). The ability to combine multiple weak learners into a strong predictor makes ensemble methods preferred for many predictive modeling applications.

Neural Networks Excel at Complex Pattern Recognition: Deep learning architectures demonstrated superior capability in capturing nonlinear interactions and complex patterns. Wide neural networks achieved 98.88% accuracy in PV fault classification with minimal electrical features, while CNN-LSTM ensembles achieved 98.7% F1-score in cybersecurity. The ability to learn hierarchical representations makes neural networks suitable for complex input spaces.

Reinforcement Learning Enables Adaptive Sequential Decision-Making: RL algorithms demonstrated strong performance in sequential decision-making tasks. SAC with ACL achieved 89% success in autonomous navigation, while Soft PPO achieved 97.7% accuracy in adaptive CAPTCHA defense. Integration of reward shaping, curriculum learning, and behavioral telemetry significantly improved performance in complex dynamic environments.

Causal Reasoning Enhances Robustness and Sample Efficiency: Integration of structural causal models with RL enables substantial improvements in sample efficiency and robustness. Key results include reduced communication complexity in federated learning ($O(d^2)$ to $O(d)$), linear speedup in sample complexity, and consistent dominance of counterfactual policies over interventional policies in environments with unobserved confounders.

Domain-Specific Considerations are Critical: Optimal model selection depends on domain-specific requirements. Cybersecurity prioritizes real-time inference and adversarial robustness. PV fault diagnosis requires high classification accuracy with computationally efficient architectures. Pharmaceutical applications demand interpretability for formulation development. Rainfall forecasting requires handling spatial and temporal dependencies with regional calibration. Autonomous navigation demands real-time decision-making with adaptability.

Challenges Remain Significant: Limited standardized benchmarks, "black box" nature of models, computational and scalability constraints, overfitting concerns, integration challenges with legacy systems, data quality issues, and difficulties in causal structure learning and counterfactual reasoning persist.

The convergence of machine learning capabilities with domain-specific knowledge has created unprecedented opportunities across critical domains. As algorithms evolve and computational resources expand, machine learning will play an increasingly central role in enabling intelligent, autonomous, and adaptive systems. The comparative framework offers guidance for algorithm selection, while identified challenges and future directions highlight opportunities for advancing the field.

The path forward requires interdisciplinary collaboration between machine learning researchers, domain experts, and practitioners to develop robust, interpretable, and deployable

models. By addressing challenges of data quality, model interpretability, computational efficiency, causal reasoning, and real-world validation, the field can realize the full potential of machine learning for transformative impact across critical domains.

References

Badr, M. M., Hamad, M. S., Abdel-Khalik, A. S., Hamdy, R. A., Ahmed, S., & Hamdan, E. (2021). Fault identification of photovoltaic array based on machine learning classifiers. *IEEE Access, 9*, 159113-159132.

Bareinboim, E., Zhang, J., & Lee, S. (2026). An introduction to causal reinforcement learning. *Causal Artificial Intelligence Lab, Columbia University*.

Bhuktar, D. M. et al. (2025). Real-time cyber threats and unauthorized access detection using embedded AI. *2025 International Conference on Computing Technologies & Data Communication (ICCTDC)*, pp. 1-5.

Breiman, L. (2001). Random forests. *Machine Learning, 45*(1), 5-32.

C S, A. et al. (2025). Artificial intelligence (AI) driven threat detection and mitigation using machine learning techniques. *2025 8th International Conference on Computing Methodologies and Communication (ICCMC)*, pp. 1629-1635.

Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the ACM SIGKDD Conference*, pp. 785-794.

Chen, Z., Han, F., Wu, L., Yu, J., Cheng, S., Lin, P., et al. (2018). Random forest based intelligent fault diagnosis for PV arrays using array voltage and string currents. *Energy Conversion and Management, 178*, 250-264.

El-Hajj, M. (2025). AI-powered threat detection and response: Leveraging machine learning for real-time intrusion detection systems (IDS) using network traffic data. *2025 5th Intelligent Cybersecurity Conference (ICSC)*, pp. 84-90.

Endalie, D., Haile, G., & Taye, W. (2021). Deep learning model for daily rainfall prediction: Case study of Jimma, Ethiopia. *Water Supply, 22*(3), 3448.

Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *Annals of Statistics, 29*(5), 1189-1232.

Fujimoto, S., van Hoof, H., & Meger, D. (2018). Addressing function approximation error in actor-critic methods. *Proceedings of the 35th International Conference on Machine Learning*, pp. 1587-1596.

Gatenbee, A., Reith, M., & Rose, A. (2026). Reviewing machine learning algorithms for threat detection in cybersecurity. *Proceedings of the 25th European Conference on Cyber Warfare & Security (ECCWS 2026)*, pp. 1051-1056.

Gemmechis, W. A. (2025). Testing machine learning algorithms for rainfall modeling. *TechRxiv*.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

Haarnoja, T., Zhou, A., Abbeel, P., & Levine, S. (2018). Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. *Proceedings of the 35th International Conference on Machine Learning*, pp. 1861-1870.

Indukuri, M., Naseerkhan, E., Rose, J., Tran, M., & Park, Y. (2026). Designing CAPTCHA systems with reinforcement learning for adaptive defense. *Department of Computer Engineering, San Jose State University*.

Kapucu, C., & Cubukcu, M. (2021). A supervised ensemble learning method for fault diagnosis in photovoltaic strings. *Energy, 227*, 120463.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*(7553), 436-444.

Leite, J. C., Nascimento, M. M. do, Almeida, A. A. S. de, Souza, M. P. de, Martins, L. da S., Silva, P. R. F. da, & Brasil, M. H. G. (2026). Development of an integrated intelligent manufacturing system utilizing IoT and data mining for the digital transformation of Industry 4.0. *ITEGAM-JETIA, 12*(59), 1059-1070.

Mak, K. K., & Pichika, M. R. (2019). Artificial intelligence in drug development: Present status and future prospects. *Drug Discovery Today, 24*(3), 773-780.

Manikandan, K. P., Reddy Onteddu, N., & Chilimi, A. K. (2025). Next-gen malware detection using AI: AI-powered malware threat detection automated malware classification through machine learning. *2025 International Conference on Intelligent Computing and Control Systems (ICICCS)*, pp. 581-588.

Mnih, V., et al. (2015). Human-level control through deep reinforcement learning. *Nature, 518*(7540), 529-533.

Pandey, P. et al. (2025). AI-powered defenses: A machine learning approaches in cybersecurity threat detection. *2025 8th International Conference on Circuit, Power & Computing Technologies (ICCPCT)*, pp. 394-399.

Pandey, R., & Patil, N. Y. (2026). A literature survey on Industry 4.0 technologies enabling real-time monitoring, predictive maintenance, and improved decision-making for enhanced efficiency and reduced operational costs. *Global Mansarovar University*.

Petrenko, D. V., & Protasov, A. G. (2026). Deep reinforcement learning-based mapless navigation algorithm for Husky robot with adaptive curriculum learning. *KPI Science News, 2*, 36-43.

Polinati, A. K. (2025). AI and deep learning-powered threat intelligence and automated response mechanisms. *2025 3rd International Conference on Sustainable Computing and Data Communication Systems (ICSCDS)*, pp. 1504-1509.

Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.

Shah, A. V., Shah, P. K., & Pandya, H. B. (2026). Comparative analysis of machine learning algorithms for predictive drug delivery systems. *International Journal of Drug Delivery Technology, 16*(15s), 190-197.

Silver, D., et al. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature, 529*(7587), 484-489.

Soori, M., & Azizi, A. (2026). Autonomous control systems using artificial intelligence, machine learning, and digital twins in Industry 4.0. *Journal of Complex Multiphysics Engineering Systems, 1*(3), 239-264.

Srilakshmi, P. et al. (2025). Real-time IoT cybersecurity using machine learning-based AI threat detection system to train generative robots. *2025 5th International Conference on Trends in Material Science and Inventive Materials (ICTMIM)*, pp. 1124-1130.

Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. MIT Press.

Taah, P. T., Asoh, D. A., Mungwe, J. N., Nguimfack, J.-D.-D., & Agoons, D. (2026). Fault diagnosis in photovoltaic systems using machine learning algorithms. *Smart Grid and Renewable Energy, 17*, 131-156.