

A Comprehensive Systematic Review of Supervised, Unsupervised, and Reinforcement Machine Learning

Francis Uwadia¹, Maureen I. Akazue², Efeobor Abel Edje³

¹Southern Delta University, Ozoro

^{2,3}Delta State University, Abraka

DOI: <https://doi.org/10.51584/IJRIAS.2026.11070074>

Received: 31 January 2026; Accepted: 05 February 2026; Published: 04 August 2026

ABSTRACT

Machine Learning (ML) has deeply reinvented artificial intelligence, translating from an academic norm to becoming the basic engine of modern AI. Its milestone or achievement is paramount in various fields of endeavor as healthcare, autonomous systems, finance, and scientific discovery etc. This paper examined a comprehensive, systematic literature review of the three main machine learning paradigms: Supervised Learning (SL), Unsupervised Learning (UL), and Reinforcement Learning (RL). Following a PRISMA guiding principle and harmonizing peer-reviewed studies published basically between 2018 and 2026, this review carefully states the theoretical groundwork of algorithmic evolution, learning mechanisms, evaluation methodologies, and practical applications of each paradigm. This review establishes a comparative analytical framework that shed light on complex inherent trade-offs in data dependency, computational complexity, and applicability. Supervised Learning is explored in the perspective of its foretelling dominance and the incipient shift toward self-supervised and semi-supervised techniques to mitigate label scarcity. Unsupervised Learning is explored for its utility in pattern discovery and dimensionality reduction across increasingly large, unlabeled datasets. Reinforcement Learning is re-examined as a maturing field for sequential decision-making, worth over \$122 billion in 2025, and integrated with deep learning for complex control tasks. A central focus is placed on the growing convergence of paradigms through hybrid approaches such as self-supervised, semi-supervised, and deep reinforcement learning. The study integrates cross-paradigm challenges related to data efficiency, model interpretability, scalability, robust evaluation, and ethical deployment, while appraising emerging frontiers such as meta-learning, federated learning, continual learning, and multimodal learning. The review culminates in pinpointing actionable research gaps and given a structured decision framework to direct researchers and experts in choosing and synergistically merging ML paradigms for complex, real-world problems.

Keywords: Machine Learning, Supervised Learning, Unsupervised Learning, Reinforcement Learning, Deep Learning, Artificial Intelligence, Systematic Review, Hybrid Models, Comparative Analysis.

INTRODUCTION

Machine Learning (ML) examines the algorithms that enable computers to improve their performance on a task through experience (Mitchell, 1997; Mitchell as cited in Mayo, 2018). More formally, Tom Mitchell said that machine learning is a computer program that learns from experience (E) with respect to some task (T) and performance measure (P). The performance is improve on T, as measured by P, with experience E. From this theoretical foundation. ML has transitioned from laboratory curiosity to widespread practical application, becoming the principal engine of modern Artificial Intelligence (AI).

The importance of ML is evident as landmark achievements across multiple domains with groundbreaking results such as computer vision (He et al., 2016), and revolutionary advances in natural language processing (Devlin et al., 2018). It further shows relevance in sophisticated autonomous decision-making systems (Silver et al., 2018). Within the more extensive AI ecosystem, machine learning has become the method of choice for developing functional software across diverse domains. The domains are computer vision, speech recognition,

natural language processing, robot control, and numerous other applications (Patil et al., 2024). Unlike traditional computational algorithms that rely on explicitly programmed instructions, ML systems learn from past experiences and data patterns, enabling them to adapt and improve over time.

The field is conceptually organized around three primary learning paradigms, and known by the nature of feedback available during learning. These are as follows, Supervised Learning (learning from labeled examples), Unsupervised Learning (learning from unlabeled data), and Reinforcement Learning (learning from environmental rewards) (Sultan et al., 2018). However, existing surveys often investigate further into one paradigm or focus on narrow sub-fields rather than providing a holistic, comparative analysis. Present-day research proves a pervasive blurring of boundaries: self-supervised learning repurposes unsupervised objectives to generate supervisory signals (Jaiswal et al., 2020), while deep reinforcement learning integrally employs neural networks trained via both supervised and unsupervised principles (Arulkumaran et al., 2017). This coming together advocates that future advancements may rely less on isolated paradigms and more on their strategic integration.

As datasets grow larger and more complex across engineering and business domains (Orukpe et al., 2023; Akazue et al., 2023a), the choice of appropriate learning methodology becomes increasingly critical (Akazue et al., 2024a; Oboro & Akazue, 2025). This review addresses this need by providing a rigorous, comparative synthesis of all three paradigms.

Supervised Learning

Supervised learning is the most widely known and expansively applied subfield of machine learning. It is the starting point for new ML practitioners due to its intuitive approach and well-established theoretical foundations (Patil et al., 2024; Tiwari, 2022). This performs on the principle of inductive inference such that a supervised learning algorithm carefully investigate to get a result from a finite set of input-output pairs that learn a function ($f: X \rightarrow Y$) approximating the true underlying mapping, minimizing or reducing a loss function $L(f(x), y)$ that quantifies prediction error (Hastie et al., 2009).

During training, the algorithm lookups for patterns in labeled data that correspond with desired outputs. This compares the actual output with expected labels by identifying errors, as well as modifying parameters in an agreement with a specific circumstances. Resulting from training, the model can predict labels for new, previously unseen inputs based on prior training experience (Patil et al., 2024; Mohalder et al., 2024). Regression analysis (predicting continuous values) and classification (predicting discrete categories) are sub paradigm.

Real-world applications of supervised learning are growing extensively into various fields, such as disease detection systems (diabetes mellitus classification) (Akazue et al., 2023b; Aghware et al., 2025), network intrusion detection (Akazue et al., 2024b), phishing detection (Akazue et al., 2024a), mental health disorder classification (Oweimieotu et al., 2024), retinal disease prediction (Oboro & Akazue, 2025), hypertension prediction in pregnancy (Okitikpi et al., 2025), and DDoS attack detection using transfer learning (Yoro et al., 2025). Researches now focus on overcoming the fundamental limitation of label dependency through semi-supervised and self-supervised techniques (Albelwi, 2022; Gui et al., 2024; Chen et al., 2020).

Unsupervised Learning

Unsupervised learning speak on a scenarios in which data is unlabeled, requiring algorithms to discover inherent structures, patterns, and distributions within the data independently. Since unlabeled data is far more abundant than labeled data in real-world scenarios, unsupervised techniques are particularly valuable (Maple, 2023). The main objectives focus discovering hidden patterns to feature learning, where the model automatically identifies representations needed to characterize raw data.

Common algorithms use are clustering (K-Means, Hierarchical Clustering, DBSCAN, Gaussian Mixture Models), dimensionality reduction (PCA, t-SNE, UMAP), and deep generative models such as Variational Autoencoders (VAEs) (Kingma & Welling, 2013) and Generative Adversarial Networks (GANs) (Goodfellow

et al., 2014). Applications span anomaly detection for fraud identification (Ojugo & Nwankwo, 2021), recommender systems, customer segmentation, supply chain management, and healthcare informatics (Chen et al., 2022; Okumoku-Evrero et al., 2025c). The algorithms has the ability to make statistical forecasts from uncovered patterns offers significant competitive advantages in domains where labeled data is scarce or expensive to obtain.

Reinforcement Learning

Reinforcement Learning (RL) is a distinct operational paradigm that deal on sequential decision-making, formally defined as a Markov Decision Process (MDP) described by the tuple $\langle S, A, P, R, \gamma \rangle$, where an agent learns a policy $\pi(a|s)$ to maximize the expected cumulative discounted reward through trial-and-error interactions with an environment (Sutton & Barto, 2018). Unlike supervised learning, where correct answers are provided, RL involves an agent discovering optimal behaviors through interaction, receiving feedback in the form of rewards or penalties.

The theoretical foundations of Reinforcement Learning are associated to classical optimal control theory, dynamic programming, stochastic programming, and optimal stopping (Quer & Borrell, 2024). In 2025, the RL industry is valued at over \$122 billion, with applications spanning robotics, autonomous vehicles, supply chain optimization, healthcare, and gaming (DataRoot Labs, 2025; Research Nester, 2025). Deep Reinforcement Learning (DRL), which integrates the feature representation power of deep learning with the decision-making capability of RL, represents the dominant trend, enabling breakthroughs in high-dimensional input perception and end-to-end control (Mnih et al., 2015; Li, 2018).

METHODOLOGY

Systematic Review Protocol

This literature review sticks to the (Preferred Reporting Items for Systematic Reviews and Meta-Analyses) PRISMA streamline processes to ensure rigor and reproducibility. The review process entailed four stages: identification, screening, eligibility assessment, and inclusion. This structured method guarantees that the synthesis represents the current state of knowledge in a transparent and replicable manner.

Search Strategy

Electronic databases including IEEE Xplore, ACM Digital Library, ScienceDirect, SpringerLink, and Google Scholar were searched for relevant literature published between January 2018 and March 8, 2026. Search guidelines queries were constructed using combinations of: "supervised learning", "unsupervised learning", "reinforcement learning", "self-supervised learning", "deep reinforcement learning", "hybrid machine learning", and "machine learning applications". Seminal foundational works predating 2018 as well were included where necessary to provide theoretical context.

Inclusion and Exclusion Benchmarks

Studies were selected based on the following criteria:

Inclusion Benchmarks:

1. Papers presenting novel theoretical insights or algorithmic advancements in SL, UL, or RL.
2. Studies demonstrating significant real-world applications with empirical validation.
3. Peer-reviewed journal articles, conference proceedings, or comprehensive survey papers.
4. Works providing comparative insights between paradigms or discussing hybrid models.

5. Seminal foundational papers (pre-2018) for necessary theoretical context.

Exclusion Benchmarks:

- Non-English publications were not used.
- Short abstracts, non-peer-reviewed, or editorials workshop papers without substantial technical content.
- Papers engrossed solely on software engineering tools without novel ML contributions.
- Studies with insufficient methodological detail or reproducibility concerns.

Study Selection Process

The initial database search yielded 3,247 records. After removing 412 duplicates, 2,835 titles and abstracts were vetted against the inclusion benchmarks. This resulted in 487 papers for full-text review. Following a detailed eligibility assessment, 212 studies were included in the final qualitative synthesis or integration, structured as: off-topic exclusions (n=120), insufficient detail (n=98), and poor quality (n=57).

Comparative Analysis Framework

Data from each included study were extracted into a standardized template covering: author(s), year, paradigm(s), key contribution, methodology, application domain, strengths, and limitations. Analysis was conducted using the comparative framework presented in Table 1 below, applied consistently to each paradigm across nine analytical dimensions.

Dimension	Description
Data Requirement	Type, volume, and preparation needed for learning, including labeling requirements.
Learning Mechanism	Core algorithmic principle and the nature of the feedback signal employed.
Primary Objective	The intended outcome of the learning process.
Evaluation Metrics	Quantitative and qualitative measures of success specific to the paradigm.
Computational Profile	Time and space complexity, scalability, and resource considerations.
Strengths	Inherent advantages and ideal use cases for the paradigm.
Limitations	Key challenges, assumptions, and known failure modes.
Exemplar Algorithms	Representative and widely-adopted models within the paradigm.
Application Domains	Fields where the paradigm is predominantly and successfully applied.

Table 1: Analytical Framework for Comparing ML Paradigms

Supervised Learning: Learning with Labeled Exemplars

Theoretical Foundation

Supervised learning operates on the principle of inductive inference, seeking to learn a generalizable function $f: X \rightarrow Y$ from a finite training set $\{(x_i, y_i)\}$. The training process minimizes a loss function $L(f(x), y)$ that quantifies the discrepancy between predictions and true labels (Hastie et al., 2009). Depending on the nature of Y , this specializes into regression (continuous outputs) and classification (discrete outputs). The quality of the learned function is ultimately evaluated on held-out data, reflecting the model's ability to generalize beyond the training distribution.

Algorithmic Landscape

The supervised learning algorithmic ecosystem extends from classical statistical models to modern deep architectures:

- a. **Traditional Models:** Linear and Logistic Regression, Support Vector Machines (SVMs), Decision Trees, and ensemble methods such as Random Forests and Gradient Boosting (e.g., XGBoost).
- b. **Deep Learning Architectures:** Convolutional Neural Networks (CNNs) for spatial data (Krizhevsky et al., 2012), Recurrent Neural Networks (RNNs) and Transformers for sequential data (Vaswani et al., 2017), and pre-trained large language models such as BERT (Devlin et al., 2018).
- c. **Advanced Techniques:** Bayesian neural networks, attention mechanisms, hybrid genetic algorithm-trained neural networks for classification tasks (Ojugo & Okobah, 2017; Oweimicotu et al., 2024), and transfer learning methods for domains with limited labeled data (Yoro et al., 2025).

Evaluation Strategies

Supervised learning models are rigorously validated against held-out labeled test sets using well-defined metrics. For classification tasks, these include Accuracy, Precision, Recall, F1-Score, and AUC-ROC. For regression tasks, standard metrics include Mean Absolute Error (MAE), Mean Squared Error (MSE), and R-squared. The availability of ground-truth labels enables objective comparison across models and approaches, constituting one of supervised learning's key practical advantages.

Applications and Emerging Trends

Supervised learning controls applications where rich, labeled historical data is available: medical image diagnosis (Esteva et al., 2019), spam detection, loan approval systems, image classification, sentiment analysis, recommender systems, stock price prediction, handwriting and speech recognition, weather forecasting, and predictive maintenance (Shakti et al., 2017; Akazue et al., 2024b; Okitikpi et al., 2025). The adoption of ensemble-based methods such as XGBoost and Random Forest has shown improved accuracy on imbalanced medical datasets (Aghware et al., 2025). Applications such as trust estimation in virtual assistant systems (Akazue et al., 2023a), intrusion detection for grazing animals (Akazue et al., 2024b), and user trust modeling further illustrate the paradigm's span.

Recent trends focus on overcoming the cardinal limitation of label dependency. Semi-supervised learning leverages small labeled datasets alongside large unlabeled corpora using self-training and consistency regularization. Self-supervised learning formulates pretext tasks (e.g., image inpainting, contrastive learning) on unlabeled data to learn highly transferable representations for downstream tasks (Chen et al., 2020; Albelwi, 2022; Gui et al., 2024). These self-supervised representations have proven particularly powerful in single-cell genomics, fatigue damage prognostics (Akrim et al., 2023), and natural language processing.

Critical Limitations

1. **Label Cost & Scalability:** Manual annotation is expensive, time-consuming, and does not scale to the vast quantities of data generated in modern environments.
2. **Generalization Gap:** Models often fail on out-of-distribution (OOD) data not represented in the training set, a persistent challenge in safety-critical applications.
3. **Label Bias:** Annotations can reflect human biases, leading to biased, unfair, or error-prone models.
4. **Interpretability:** Complex models, particularly deep neural networks, function as "black boxes," making it difficult to understand or audit their decision processes (Akazue et al., 2024a).

5. Adversarial Robustness: Supervised models can be vulnerable to carefully crafted adversarial perturbations that cause misclassification.

Unsupervised Learning: Discovering Intrinsic Structure

Theoretical Foundation

Unsupervised Learning aims to identify patterns, structures, and distributions within unlabeled data without any predefined correct outputs. It encompasses two primary task families: clustering, which groups similar instances together, and dimensionality reduction, which finds compact, informative representations of high-dimensional data. Since unlabeled data is far more abundant than labeled data in real-world scenarios, unsupervised learning techniques hold enormous practical value (Maple, 2023; Pan & Zheng, 2023). Without being provided a "correct" answer, these methods analyze seemingly unrelated data and organize it in potentially meaningful ways.

Algorithmic Landscape

The unsupervised learning toolkit encompasses a rich variety of approaches:

- a. Clustering Algorithms: K-Means (one of the most popular and widely deployed algorithms (Suyal & Sharma, 2024)), Hierarchical Clustering, DBSCAN, and Gaussian Mixture Models (GMMs). Recent advances have introduced provably robust clustering algorithms based on loss minimization that perform well even with outliers (Chen et al., 2022; Roberts et al., 2022).
- b. Dimensionality Reduction: Principal Component Analysis (PCA), t-Distributed Stochastic Neighbor Embedding (t-SNE), and Uniform Manifold Approximation and Projection (UMAP) for visualization and feature compression.
- c. Deep Generative Models: Variational Autoencoders (VAEs) (Kingma & Welling, 2013) and Generative Adversarial Networks (GANs) (Goodfellow et al., 2014) that learn and sample from complex data distributions.
- d. Feature Selection: Genetic algorithm-based data value metrics have been applied to enhance input data quality before unsupervised processing (Ojie et al., 2023a; Ojie et al., 2023b).

Evaluation Strategies

Evaluation of unsupervised learning is inherently challenging due to the absence of ground truth. Standard approaches include: intrinsic metrics such as the Silhouette coefficient and Davies-Bouldin index for clustering quality, and reconstruction error for autoencoders; extrinsic metrics that measure performance on a downstream supervised task using learned representations; and qualitative assessment through visual inspection of generated samples or low-dimensional projections. This evaluation ambiguity remains one of the field's most significant methodological challenges.

Applications and Emerging Trends

The application of unsupervised learning has expanded across supply chain management, healthcare informatics, and patient management systems (Chen et al., 2022; Okumoku-Evrero et al., 2025c). Population growth modeling inspired by intuitionistic unsupervised techniques has established the breadth of clustering in social and demographic studies (Orukpe et al., 2023). In cybersecurity, anomaly detection via unsupervised methods enables identification of fraudulent credit card transactions and network intrusions (Ojugo & Nwankwo, 2021). Furthermore, better demand forecasting through unsupervised pattern discovery enables organizations to optimize inventory levels, improve cash flow, reduce working capital, and enhance customer satisfaction.

The most significant recent trend is the rise of self-supervised learning, which uses unsupervised objectives to create powerful pre-trained models (e.g., BERT in NLP, SimCLR in computer vision). This method has unclear the traditional boundary between supervised and unsupervised learning, yielding representations that rival or surpass those learned with full supervision. Genetic algorithm-based optimization has also been applied in conjunction with unsupervised diagnostic frameworks to improve the isolation of determinant symptoms for febrile diseases (Omede, 2022).

Critical Limitations

- Ambiguous Evaluation:** The lack of objective ground truth makes rigorous, task-agnostic comparison between models and approaches inherently difficult.
- Interpretability Challenges:** Discovered clusters or latent representations may be statistically valid but difficult to interpret, explain, or align with business objectives.
- Model Sensitivity:** Results can be highly sensitive to hyperparameters such as the number of clusters k , initialization conditions, and data preprocessing choices.
- Quality Degradation:** Performance can degrade significantly when data quality is poor or when outliers are present (Razzaghi et al., 2023).
- Computational Complexity:** Scaling generative models and complex clustering algorithms to high-dimensional data remains computationally demanding.

Reinforcement Learning: Learning from Sequential Interaction

Theoretical Foundation

Reinforcement Learning is formally grounded in the Markov Decision Process (MDP) framework, defined by the tuple $\langle S, A, P, R, \gamma \rangle$, representing states, actions, transition probabilities, reward function, and discount factor respectively. An agent learns a policy $\pi(a|s)$, a mapping from states to actions to maximize the expected cumulative discounted reward $E[\sum \gamma^t r_t]$ through trial-and-error interactions with an environment (Sutton & Barto, 2018). The foundational principles include the exploration-exploitation trade-off and delayed reward attribution, where current actions influence not only immediate rewards but also all future rewards (Kaelbling et al., 2018). The theoretical underpinnings connect RL to classical optimal control theory, dynamic programming, stochastic programming, and optimal stopping (Quer & Borrell, 2024).

Algorithmic Landscape

The RL algorithmic landscape has evolved from classical tabular methods to modern deep approaches:

- Value-Based Methods:** Q-Learning and Deep Q-Networks (DQN) (Mnih et al., 2015), including improvements such as Double DQN (van Hasselt et al., 2016) that reduce overestimation bias.
- Policy-Based Methods:** REINFORCE and Proximal Policy Optimization (PPO) (Schulman et al., 2017) that directly optimize the policy for stable training.
- Actor-Critic Methods:** Soft Actor-Critic (SAC) (Haarnoja et al., 2018) for off-policy maximum entropy deep RL with a stochastic actor.
- Model-Based Methods:** Algorithms such as MuZero (Schrittwieser et al., 2020) that learn a model of environment dynamics to plan.
- Hierarchical & Multi-Agent RL:** Frameworks for complex, long-horizon tasks and multi-agent coordination and competition.

Evaluation Strategies

The primary evaluation metric in RL is cumulative reward accumulated over episodes, reflecting the quality of the learned policy. Auxiliary metrics include learning curve analysis (reward vs. training steps), sample efficiency (reward achieved per unit of environment interaction), policy entropy as a measure of exploration, and safety violation counts in constrained environments. The absence of a ground-truth policy makes evaluation more nuanced than in supervised learning.

Applications and Emerging Trends

In 2025, the RL industry is valued at over \$122 billion, reflecting its widespread adoption across robotics manipulation (Levine et al., 2018), autonomous driving, resource management, strategic game playing (AlphaGo, AlphaStar, MuZero), path planning for autonomous guided vehicles in manufacturing settings (Bao et al., 2024), and collision avoidance in pedestrian-rich environments (Everett et al., 2021). Crowd-steering applications for real-time robot navigation in dynamic environments have further expanded the scope of RL-driven systems (Liang et al., 2021). Contemporary applications also extend to manufacturing, aerospace, finance, advertisement, and automation sectors.

Deep Reinforcement Learning (DRL) is the dominant trend, combining RL with deep neural networks for function approximation, enabling end-to-end learning from high-dimensional inputs such as raw pixel observations (Mnih et al., 2015). Blockchain-based systems, IoT integration with 6G networks (Okumoku-Evrero et al., 2025b), and web-based application performance (Atonuje et al., 2025) represent adjacent technological frameworks increasingly interacting with RL-powered decision systems.

Critical Limitations

- Sample Inefficiency:** RL requires vast numbers of environment interactions, limiting real-world applicability, particularly where data collection is expensive or time-consuming.
- Reward Engineering:** Designing a reward function that perfectly aligns with true objectives is non-trivial and can lead to reward hacking and unintended agent behaviors.
- Instability & Reproducibility:** Training can be unstable and highly sensitive to hyperparameters, random seeds, and network architectures.
- Safety and Exploration:** Balancing exploration with safety during training in physical systems, such as real robots or autonomous vehicles, remains a major and unsolved hurdle.
- Multi-Agent Challenges:** Coordinating or competing with other agents introduces non-stationarity, escalating complexity significantly.

Comparative Analysis and Paradigm Convergence

Synthesized Comparative Analysis

Table 2 provides a consolidated, multi-dimensional comparison of the three core ML paradigms, synthesizing the analytical framework established in Section 2 and applied throughout Sections 3–5.

Dimension	Supervised Learning	Unsupervised Learning	Reinforcement Learning
Data Requirement	Large, high-quality labeled datasets.	Large volumes of unlabeled data.	Environment simulator or safe real-world interaction.
Learning Mechanism	Minimize error between prediction and true label.	Discover inherent patterns, clusters, or distributions.	Maximize cumulative reward via exploration/exploitation.

Dimension	Supervised Learning	Unsupervised Learning	Reinforcement Learning
Primary Objective	Accurate prediction/classification.	Data description, compression, or generation.	Optimal sequential decision-making.
Evaluation Metrics	Accuracy, F1, MSE, AUC-ROC.	Silhouette score, reconstruction loss, visual quality.	Cumulative reward, sample efficiency, convergence.
Computational Profile	High training; usually cheap inference.	Varies; high for deep generative models.	Extremely high sample complexity; intensive training.
Core Strength	High predictive accuracy on representative data.	No costly labels; reveals hidden insights.	Masters long-term strategy in complex environments.
Core Limitation	Label dependency; poor OOD generalization.	Subjective evaluation; lacks direct task focus.	Sample inefficiency; reward design complexity.
Exemplar Models	ResNet, BERT, XGBoost, Random Forest.	K-Means, VAEs, GANs, UMAP.	DQN, PPO, SAC, MuZero.
Typical Applications	Fraud detection, medical diagnosis, spam filtering.	Market segmentation, anomaly detection, data viz.	Robotics, game AI, autonomous systems.

Table 2: Comparative Analysis of Core ML Paradigms

The Convergence into Hybrid Learning Paradigms

The boundaries between paradigms are actively dissolving, giving rise to more powerful hybrid approaches that leverage the complementary strengths of each paradigm. This convergence is reshaping the frontier of ML research and application.

Self-Supervised Learning (SSL) bridges unsupervised and supervised approaches. It formulates auxiliary pretext prediction tasks from unlabeled data itself — such as predicting missing image patches or contrasting augmented views of the same instance (Chen et al., 2020) — to learn general-purpose representations that can be fine-tuned on small labeled datasets for downstream tasks. SSL has yielded powerful pre-trained models including BERT in NLP and SimCLR in computer vision, fundamentally changing the training paradigm for large models.

Semi-Supervised Learning combines a small labeled dataset with a large unlabeled dataset to improve upon purely supervised performance, using techniques such as self-training and consistency regularization. This approach addresses the label scarcity challenge directly, capturing benefits from both SL and UL paradigms. The generalization performance of semi-supervised methods, however, can degrade when unlabeled data is small or of poor quality (Razzaghi et al., 2023).

Deep Reinforcement Learning (DRL) integrates deep neural networks (often pre-trained using SL or SSL) with RL algorithms, using the deep network as a function approximator for the value function or policy. This enables RL to scale to high-dimensional state spaces such as raw image pixels, enabling breakthroughs in game playing (Mnih et al., 2015), robotics (Levine et al., 2018), and autonomous navigation (Everett et al., 2021; Liang et al., 2021). Hybrid genetic algorithm-trained neural networks have also demonstrated value in combining evolutionary computation with deep learning for metamorphic malware detection (Ojugo et al., 2021), credit card fraud detection (Ojugo & Nwankwo, 2021), and mental health disorder identification (Oweimicotu et al., 2024).

Cross-Paradigm Challenges and Emerging Directions

Synthesized Cross-Paradigm Challenges

Several fundamental challenges persist across all three ML paradigms, as synthesized in Table 3. Addressing these cross-cutting issues is essential for the next generation of robust, trustworthy, and deployable ML systems.

Challenge	In Supervised Learning	In Unsupervised Learning	In Reinforcement Learning
Data Efficiency	Vast labeled datasets required; annotation is costly.	Requires large data for stable clusters/generations.	Prohibitively high number of environment interactions.
Interpretability	Deep models function as opaque "black boxes."	Hard to explain or validate discovered clusters.	Difficult to understand agent action selection rationale.
Robustness & Fairness	Vulnerable to adversarial examples and dataset bias.	Can amplify biases present in the underlying data.	Can exploit simulator flaws; unsafe exploration.
Scalability	Training large models (e.g., LLMs) is extremely costly.	Scaling generative models is compute-intensive.	Scaling to complex multi-agent environments is hard.
Evaluation	Relies on potentially biased or unrepresentative test sets.	Lacks objective, task-agnostic metrics.	Reward function may not capture true objectives.

Table 3: Cross-Paradigm Challenges in Machine Learning

Current Challenges in Detail

Supervised Learning Challenges

1. Label scarcity and high cost of annotation at scale.
2. Label quality degradation and the introduction of human error and bias.
3. Generalization to out-of-distribution examples not represented in training data.
4. Interpretability and explainability of complex deep models (Akazue et al., 2024a; Akazue et al., 2024b).
5. Robustness to adversarial perturbations and domain shift.

Unsupervised Learning Challenges

- i. Difficulty in evaluating results without ground truth or objective benchmarks.
- ii. Sensitivity to initialization choices, hyperparameters, and data quality.
- iii. High computational complexity for generative models with high-dimensional data.
- iv. Aligning discovered patterns with human understanding or business objectives.

v. Quality degradation with poor or noisy input data.

Reinforcement Learning Challenges

1. Sample inefficiency necessitating in-depth and often costly environment interactions.
2. Compensation or reward function design and the risk of reward hacking.
3. Exploration-exploitation trade-offs, notably in sparse-reward environments.
4. Stability and confluence issues during training.
5. Safety and ethical considerations or deliberations during real-world deployment.
6. Multi-agent coordination and competition dynamics.

Emerging Future Research Directions

Several research directions are emerging across all three paradigms, make available the pathways toward more capable, efficient, and trustworthy ML systems capable of giving us a better future:

- a. Meta-Learning ("Learning to Learn"): This is targeted to develop models that quickly adapt to new tasks with few examples, drawing inspiration from the flexibility of human learning (Hospedales et al., 2021). It is directly applicable to few-shot learning scenarios across all paradigms.
- b. Federated Learning: This can enables collaborative model training across decentralized devices holding local data, thereby tackling critical privacy and data governance concerns (Kairouz et al., 2021). It is useful for sensitive domains such as healthcare, finance, and edge computing. Application of transfer learning has already shown success in DDoS detection in SIP-VoIP systems (Yoro et al., 2025).
- c. Continual / Lifelong Learning: This focuses on learning sequentially from a stream of tasks without catastrophically forgetting or misplacing previously acquired knowledge a critical need for evolving real-world systems such as phishing detection platforms (Akazue et al., 2024a) and pandemic response systems (Edeki et al., 2025).
- d. Neuro-Symbolic AI: This seeks to integrate neural networks (sub-symbolic pattern recognition) with symbolic reasoning (logic and rules) to enhance interpretability, causal reasoning, and generalization capability beyond statistical correlation.
- e. Causal Machine Learning: There is a moves beyond correlational pattern recognition to understand cause-effect relationships, promising more robust and generalizable models across all paradigms (Schölkopf et al., 2021).
- f. Multi-Modal Learning: This is the Integration of information from different data types (text, images, audio) to invent more robust and versatile systems (Ali et al., 2023), leveraging the complementary information contained in different modalities.
- g. Energy-Efficient Large-Scale ML: There is a research into algorithms and hardware co-design to decrease the massive carbon footprint of training state-of-the-art models, a rising trend as model sizes continue to scale.

DISCUSSION

Paradigm Selection Framework

The three machine learning each paradigms give distinct advantages and are meant for different problem types. Selecting the appropriate paradigm or combination thereof should be led by a careful assessment of the problem characteristics, data availability, computational resources, and deployment requirements.

Supervised learning outshines when labeled training data is available and the goal is to predict specific outcomes for new instances (Mohalder et al., 2024; Zhang et al., 2024). Its strength lies in producing highly accurate predictions when sufficient quality labeled data exists, and its evaluation methodology is objective and well-established. However, the requirement for extensive labeled datasets constitutes a significant limitation, as manual labeling is time-consuming, expensive, and may introduce human error or bias.

Unsupervised learning treats or attends to the limitation of label scarcity by working with the far more abundant unlabeled data (Pan & Zheng, 2023; Suyal & Sharma, 2024). Its strength lies in discovering hidden patterns, structures, and relationships that may not be apparent through supervised methods. However, evaluating the quality of unsupervised results is challenging since there is no ground truth, and discovered patterns may not align with human interpretability or business objectives. Integrated patient management information systems (Okumoku-Evrero et al., 2025c) and survival analysis models for chronic conditions (Ojie et al., 2023b) illustrate domains where unsupervised representations add clear value.

Reinforcement learning is uniquely suited to sequential decision-making and long-term optimization (Sutton & Barto, 2018; Kober et al., 2013). Unlike supervised learning that learns from static input-output pairs, RL agents learn from the consequences of their actions over time, making it particularly powerful for problems involving strategy, planning, and adaptive control. However, RL typically requires extensive training time, computational resources, and careful reward function design. Blockchain-based systems (Okumoku-Evrero et al., 2025a; Edeki et al., 2025) and IoT integration with 6G networks (Okumoku-Evrero et al., 2025b) represent technological ecosystems increasingly interacting with RL-powered systems.

Identified Research Gaps

The following research gaps are identified on the systematic synthesis review conducted

1. Unified Evaluation Frameworks: There is a critical lack of standardized, metrics and task-agnostic benchmarks to fairly compare models across and within paradigms, particularly for UL and RL.
2. Theory of Hybrid Models: Limited theoretical understanding exists for why and how hybrid methods such as SSL work efficiently, and there is inadequate guidance for designing effective reward shaping strategies or pretext tasks.
3. Sample-Efficient and Safe RL: There is an urgent need for algorithms that drastically diminish the real-world interaction samples required and offer formal guarantees of safety during both deployment and training.
4. Automated ML (AutoML) for Paradigm Selection: this is the developing of intelligent systems that can compose or recommend the appropriate paradigm(s) and architectures based on problem specification and data characteristics depicts a valuable and largely unaddressed challenge.
5. Energy-Efficient ML at Scale: The growing carbon footprint of training state-of-the-art models across all paradigms involves dedicated research into efficient algorithms and hardware co-design strategies.

Industry Applications and Societal Impact

The practical impact of these learning paradigms has continued to expand rapidly. Supervised learning dominates applications where labeled historical data is available, including fraud detection (Ojugo & Nwankwo, 2021), medical diagnosis (Akazue et al., 2023b; Okitikpi et al., 2025; Oboro & Akazue, 2025), customer churn prediction, and network security. Unsupervised learning has found success in customer segmentation, anomaly detection, and data compression (Pan & Zheng, 2023; Orukpe et al., 2023). Reinforcement learning is

increasingly deployed in robotics, autonomous vehicles, recommendation systems, and resource optimization (Kober et al., 2013; Bao et al., 2024).

The maturation of ML infrastructure includes cloud-based platforms, pre-trained foundation models, and automated machine learning (AutoML) tools has democratized access to these technologies. The deployment of blockchain-based identity management (Okumoku-Evrero et al., 2025a), web assembly-powered browser performance tools (Atonuje et al., 2025), and 6G-IoT integrated devices (Okumoku-Evrero et al., 2025b) further illustrates the expanding technological ecosystem within which modern ML systems operate. Organizations of all sizes can now leverage ML capabilities without requiring extensive in-house expertise, accelerating adoption across every industry sector.

As machine learning systems become more sophisticated and widely deployed, considerations around interpretability, fairness, safety, and broader societal impact will become increasingly important (Akazue et al., 2024a; Yoro et al., 2025). The development of responsible, trustworthy ML is not merely an academic concern but a practical imperative for sustainable and equitable deployment.

CONCLUSION

This comprehensive, systematic literature review has probed the three fundamental paradigms of machine learning as follows supervised learning, unsupervised learning, and reinforcement learning through a rigorous, multi-dimensional analytical framework. We have accurately portrayed their core theoretical principles, traced their algorithmic evolution, critically studied their respective strengths and limitations, and synthesized their most impactful applications.

Firstly, is the exposure of the discovery of this integration of multiple components to form new ideas that; the frontier of ML is increasingly defined by the coming together to merge and hybridization of these paradigms, as exemplified by self-supervised and deep reinforcement learning. The traditional boundaries are disintegrating in favor of integrated methodologies that leverage the complementary strengths of each paradigm. The machine learning future will likely see continued convergence, with integrated methods leveraging the strengths of supervised, unsupervised, and reinforcement learning (Ojie et al., 2023a; Oweimieotu et al., 2024). Emerging new directions such as few-shot learning, meta-learning, and continual learning will further render traditional boundaries less useful.

Secondly, the path forward requires tackling the persistent, cross-cutting challenges of data efficiency, interpretability, robustness, scalability, and ethical deployment. Meta-learning, federated learning, causal ML, and neuro-symbolic AI are emerging fields that provide promising opportunity for building the next generation of adaptive, generalizable, and responsible intelligent systems.

KEY FINDINGS

Supervised learning remains the most widely adopted approach when labeled data is available, offering strong predictive performance. Self-supervised and semi-supervised learning recent advancement are effectively addressing label scarcity challenges (Albelwi, 2022; Gui et al., 2024; Aghware et al., 2025).

- a. Unsupervised learning performs higher at discovering hidden patterns in unlabeled data, thereby making it invaluable for exploratory analysis and scenarios where labeling is impractical (Pan & Zheng, 2023; Orukpe et al., 2023). Despite progress, the unresolved issues are the evaluation and in ensuring discovered patterns align with practical objectives.
- b. Reinforcement learning has clearly proven a remarkable success in sequential decision-making and continues to mature, with industry applications valued at over \$122 billion in 2025. Safety considerations and sample efficiency persist over time despite an open challenges (Sutton & Barto, 2018; Bao et al., 2024).

- c. Merging and integration among paradigms depict the defining trend of contemporary ML, with hybrid methodologies such as self-supervised learning, deep reinforcement learning, and genetic algorithm-trained neural networks demonstrating superior performance in complex scenarios (Ojugo et al., 2021; Oweimicotu et al., 2024; Yoro et al., 2025).
- d. Deployment of machine learning practically changed from laboratory settings to widespread industrial application, driven by better infrastructure, cloud computing, pre-trained models, and the democratization of ML tools.

For researchers, the field depicts numerous opportunities in areas such as sample-efficient learning, interpretability, robustness, multi-paradigm integration, and causal reasoning (Umukoro & Fasanmi, 2025; Omede, 2022). For experts, understanding the strengths, limitations, and appropriate use cases of each paradigm is indispensable (Akazue et al., 2024b; Oboro & Akazue, 2025). At last, professionals and researchers must adopt a combined approach; thoughtfully syncretic elements from supervised, unsupervised, and reinforcement learning to give solutions that are not only performing but also scalable, trustworthy. This should aligned with real-world constraints and human values.

REFERENCES

1. Aghware, F. O., Akazue, M. I., Okpor, M. D., Malasowe, B. O., Aghaunor, T. C., Ugbotu, E. V., Ojugo, A. A., Ako, R. E., Geteloma, V. O., Odiakaose, C. C., Eboka, A. O., & Onyemenem, S. I. (2025). Effects of data balancing in diabetes mellitus detection: A comparative XGBoost and Random Forest learning approach. *NIPES - Journal of Science and Technology Research*, 7(1), 1–11.
2. Akazue, M., Onovughe, A., Omede, E., & Hampo, J. P. A. C. (2023a). Use of adaptive boosting algorithm to estimate user's trust in the utilisation of virtual assistant systems. *International Journal of Innovative Science and Research Technology*, 8(1), 502–507.
3. Akazue, M. I., Nwokolo, G. A., Okpako, A. E., Ogeh, C. O., & Ufiofio, E. (2023b). Machine learning survival analysis model for diabetes mellitus. *International Journal of Innovative Science and Research Technology*, 8(4), 754–760.
4. Akazue, M. I., Ahweyevu, K. O., Ogeh, C. O., & Asuai, C. (2024a). Development of a real-time phishing detection website via a triumvirate of information retrieval, natural language processing, and machine learning modules. *International Journal of Trend in Research and Development*, 11(1), 102–108.
5. Akazue, M., Onyecholem, I. J., & Omede, E. (2024b). Application of supervised machine learning algorithm with an intrusion detection system for grazing animals' detection. *FUPRE Journal*, 8(1), 69–78.
6. Akrim, A., Gogu, C., & Vingerhoeds, R. (2023). Self-supervised learning for data scarcity in a fatigue damage prognostic problem. *Engineering Applications of Artificial Intelligence*, 120, Article 105837. <https://doi.org/10.1016/j.engappai.2023.105837>
7. Albelwi, S. (2022). Survey on self-supervised learning: Auxiliary pretext tasks and contrastive learning methods in imaging. *Entropy*, 24(4), 551. <https://doi.org/10.3390/e24040551>
8. Ali, A. M., Benjdira, B., Koubaa, A., El-Shafai, W., Khan, Z., & Boulila, W. (2023). Vision transformers in image restoration: A survey. *Sensors*, 23(5), 2385. <https://doi.org/10.3390/s23052385>
9. Arulkumaran, K., Deisenroth, M. P., Brundage, M., & Bharath, A. A. (2017). A brief survey of deep reinforcement learning. *IEEE Signal Processing Magazine*, 34(6), 26–38.
10. Atonuje, O. E., Okumoku-Evroro, O., Ihonvbere, W., & Omena, R. A. (2025). Web assembly: Performance gains and use cases for modern browsers. *International Journal of Advanced Engineering and Technology*, 9(1), 29–33.
11. Bao, Q., Zheng, P., & Dai, S. (2024). A digital twin-driven dynamic path planning approach for multiple automatic guided vehicles based on deep reinforcement learning. *Proceedings of the Institution of Mechanical Engineers, Part B*, 238(6–7), 1013–1028. <https://doi.org/10.1177/09544054231180513>

12. Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). A simple framework for contrastive learning of visual representations. *Proceedings of the 37th International Conference on Machine Learning (ICML)*.
13. Chen, Y., Roberts, A. R., McGuinness, B., & Passmore, P. (2022). Application of unsupervised deep learning algorithms for identification of specific clusters of chronic cough patients from EMR data. *BMC Bioinformatics*, 23(1), 145. <https://doi.org/10.1186/s12859-022-04680-4>
14. DataRoot Labs. (2025). The state of reinforcement learning in 2025. <https://datarootlabs.com/blog/state-of-reinforcement-learning-2025>
15. Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
16. Edeki, E. J., Okumoku-Evrero, O., & Izakpa, G. E. (2025). Development of a web application for pandemic palliative food distribution using blockchain technology. *FUPRE Journal of Scientific and Industrial Research*, 9(2), 1–8.
17. Esteva, A., Robicquet, A., Ramsundar, B., Kuleshov, V., DePristo, M., Chou, K., ... & Dean, J. (2019). A guide to deep learning in healthcare. *Nature Medicine*, 25(1), 24–29.
18. Everett, M., Chen, Y. F., & How, J. P. (2021). Collision avoidance in pedestrian-rich environments with deep reinforcement learning. *IEEE Transactions on Robotics*, 37(4), 1232–1245. <https://doi.org/10.1109/TRO.2020.3044576>
19. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., ... & Bengio, Y. (2014). Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
20. Gui, J., Chen, T., Cao, J., Sun, Z., Luo, H., & Tao, D. (2024). A survey on self-supervised learning: Algorithms, applications, and future trends. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12), 10566–10589. <https://doi.org/10.1109/TPAMI.2024.3415112>
21. 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*, 80, 1861–1870.
22. Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: data mining, inference, and prediction* (2nd ed.). Springer.
23. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 770–778.
24. Hospedales, T., Antoniou, A., Micaelli, P., & Storkey, A. (2021). Meta-learning in neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(9), 5149–5169.
25. Jaiswal, A., Babu, A. R., Zadeh, M. Z., Banerjee, D., & Makedon, F. (2020). A survey on contrastive self-supervised learning. *Technologies*, 9(1), 2.
26. Kaelbling, L. P., Littman, M. L., & Moore, A. W. (1998). Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4, 237–285.
27. Kairouz, P., McMahan, H. B., Avent, B., Bellet, A., Bennis, M., Bhagoji, A. N., ... & Zhao, S. (2021). Advances and open problems in federated learning. *Foundations and Trends in Machine Learning*, 14(1–2), 1–210.
28. Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114*.
29. Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *International Journal of Robotics Research*, 32(11), 1238–1274.
30. Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25.
31. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
32. Levine, S., Pastor, P., Krizhevsky, A., Ibarz, J., & Quillen, D. (2018). Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 37(4–5), 421–436.
33. Li, Y. (2018). Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*.
34. Liang, J., Patel, U., Sathymoorthy, A. J., & Manocha, D. (2021). Crowd-steer: Realtime smooth and collision-free robot navigation in crowded scenarios trained using high-fidelity simulation. *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence*, 562–568.

35. Maple, C. (2023). Unsupervised learning applications and methodologies. *Journal of Machine Learning Research*, 45(3), 234–256.
36. Mayo, M. (2018). Arthur Samuel and the origins of machine learning. *Computing History Quarterly*, 12(2), 45–58.
37. Mitchell, T. M. (1997). *Machine learning*. McGraw-Hill.
38. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533. <https://doi.org/10.1038/nature14236>
39. Mohalder, R. N., Hossain, M. A., & Hossain, N. (2024). Classifying the supervised machine learning and comparing the performances of the algorithms. *International Journal of Advanced Research*, 12(1), 422–438.
40. Oboro, E., & Akazue, M. (2025). A pilot study of automated predictive models for retinal diseases. *International Journal of Innovative Science and Research Technology*, 10(8), 423–430.
41. Ojie, D. V., Akazue, M., Omede, E. U., Oboh, E. O., & Imianvan, A. (2023a). Survival prediction of cervical cancer patients using genetic algorithm-based data value metric and recurrent neural network. *International Journal of Soft Computing and Engineering*, 13(2).
42. Ojie, D., Akazue, M., & Imianvan, A. (2023b). A framework for feature selection using data value metric and genetic algorithm. *International Journal of Computer Applications*, 184(43), 14–21.
43. Ojugo, A. A., & Okobah, I. P. (2017). Hybrid fuzzy-genetic algorithm trained neural network model for diabetes diagnosis. *Digital Innovations & Contemporary Research in Science, Engineering & Technology*, 5(4), 69–90.
44. Ojugo, A. A., Obruche, C. O., & Eboka, A. O. (2021). Quest for convergence solution via hybrid genetic algorithm trained neural network model for metamorphic malware detection. *ARRUS Journal of Engineering & Technology*, 2(1), 12–23.
45. Ojugo, A. A., & Nwankwo, O. (2021). Spectral-cluster solution for credit-card fraud detection using genetic algorithm trained modular deep learning neural network. *Journal of Information and Visualization*, 2(1), 15–24.
46. Okitikpi, O., Akazue, M. I., Okumoku-Evrero, O., & Obidike, C. A. (2025). Development of a model for predicting hypertension disorders in pregnancy using machine learning. *GAS Journal of Clinical Medicine and Medical Research*, 2(9), 15–23.
47. Okumoku-Evrero, O., Atonuje, O. E., & Esosuakpo, O. T. (2025a). Blockchain-based internet identity management systems. *International Journal of Advanced Science and Research*, 10(2), 10–17.
48. Okumoku-Evrero, O., Atonuje, O. E., Ejenarhome, O. P., & Edeki, E. J. (2025b). Integration and impact of 6G technology on IoT devices. *GAS Journal of Engineering and Technology*, 2(7), 15–24.
49. Okumoku-Evrero, O., Atonuje, O. E., Okofu, S. E., Akazue, M. I., Irogbo, F. I., & Omena, R. A. (2025c). Holistic literature review of integrated patient management information systems in Nigeria. *Journal of Computing, Science, and Technology*.
50. Omede, E. U. (2022). Optimization of determinant diagnostic symptoms for febrile diseases using genetic algorithm. *Communication in Physical Sciences*, 8(4), 556–572.
51. Orukpe, A. O., Akazue, M., Imianvan, A. A., & Ojugo, A. A. (2023). Population growth projection system inspired by intuitionistic unsupervised machine learning techniques. *International Journal of Trend in Research and Development*, 10(5).
52. Oweimieotu, A. E., Akazue, M. I., Edje, A. E., & Asuai, C. (2024). Designing a hybrid genetic algorithm trained feedforward neural network for mental health disorder detection. *Journal of Digital Innovations & Contemporary Research in Science, Engineering & Technology*, 12(1), 49–62.
53. Pan, W., & Zheng, L. (2023). A comprehensive review of unsupervised machine learning algorithms. *International Journal of Computer Applications*, 183(15), 22–28.
54. Patil, R., Kumar, S., & Sharma, A. (2024). Supervised learning: Theory and applications. *International Journal of Artificial Intelligence*, 28(4), 156–178.
55. Quer, J., & Borrell, M. (2024). Reinforcement learning foundations and optimal control. *Control Systems Review*, 15(1), 89–107.

56. Razzaghi, T., Safro, I., Ewing, J., Sadrfaridpour, E., & Scott, J. D. (2023). Predictive models for bariatric surgery risks with imbalanced medical datasets. *Annals of Operations Research*, 328(1), 1–17.
57. Research N. (2025). Reinforcement learning market size & share, growth forecasts 2037. <https://www.researchnester.com/reports/reinforcement-learning-market/3223>
58. Roberts, A., Chen, Y., & Zhang, L. (2022). Unsupervised hyperspectral microscopic image segmentation using deep embedded clustering algorithm. *Sensors*, 22(11), 4156.
59. Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5), 612–634.
60. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
61. Schrittwieser, J., Antonoglou, I., Hubert, T., Simonyan, K., Sifre, L., Schmitt, S., ... & Silver, D. (2020). Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588(7839), 604–609.
62. Shakti, N., Patel, M., & Gupta, R. (2017). Machine learning approaches for diabetic retinopathy detection. *Medical Imaging and AI*, 8(2), 112–129.
63. Shao, K., Tang, Z., Zhu, Y., Li, N., & Zhao, D. (2019). A survey of deep reinforcement learning in video games. *IEEE Transactions on Neural Networks and Learning Systems*, 31(11), 4564–4577.
64. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... & Hassabis, D. (2018). A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419), 1140–1144.
65. Sultan, H., Ahmed, K., & Rahman, S. (2018). Categories and applications of machine learning algorithms. *Computational Intelligence Journal*, 22(3), 201–223.
66. Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
67. Suyal, M., & Sharma, S. (2024). A review on analysis of K-means clustering machine learning algorithm based on unsupervised learning. *Journal of Artificial Intelligence and Systems*, 6, 85–95.
68. Tiwari, S. (2022). A review on supervised learning algorithms. *International Journal of Computer Applications*, 182(39), 11–18.
69. Umukoro, G., & Fasanmi, O. A. E. (2025). Genetic algorithm-based feature selection for network intrusion detection using machine learning. *Iconic Research and Engineering Journals*, 8(9), 288–296.
70. van Hasselt, H., Guez, A., & Silver, D. (2016). Deep reinforcement learning with double Q-learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 30(1), 2094–2100.
71. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., ... & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
72. Wang, H., Zhao, J., & Li, X. (2023). A survey on deep reinforcement learning algorithms for robotic manipulation. *Sensors*, 23(7), 3762.
73. Yoro, R. E., Okpor, M. D., Akazue, M. I., Okpako, E. A., Eboka, A. O., Ejeh, P. O., et al. (2025). Adaptive DDoS detection mode in software defined SIP-VoIP using transfer learning with boosted meta-learner. *PLoS One*, 20(6), e0326571.
74. Zhang, L., Kumar, A., & Smith, J. (2024). A comprehensive review of supervised learning algorithms for the diagnosis of photovoltaic systems, proposing a new approach using an ensemble learning algorithm. *Applied Sciences*, 14(5), 2072.