

Systematic Review

Applications of Machine Learning Across Smart Manufacturing, Healthcare, Finance, Computer Vision, Robotics, and Environmental & Sustainability: A Systematic Literature Review

Narjes Sadeghiamirshahidi * , Seyedeh Elham Kamali and Bhavani Rath Reddy Dere

Department of Mechanical and Industrial Engineering, University of New Haven, 300 Boston Post Rd, West Haven, CT 06516, USA; skama12@unh.newhaven.edu (S.E.K.); bdere3@unh.newhaven.edu (B.R.R.D.)

* Correspondence: nsadeghi@newhaven.edu

Featured Application

This study provides a structured, cross-domain framework for selecting and deploying machine learning (ML) models in real-world systems, supporting data-driven decision-making across smart manufacturing, healthcare, finance, robotics, and environmental sustainability. By synthesizing methodological trends, deployment challenges, and domain-specific requirements, the review enables practitioners and policymakers to identify appropriate ML paradigms based on data availability, interpretability needs, and operational constraints. The findings facilitate the design of robust, explainable, and scalable ML solutions for applications such as predictive maintenance, clinical decision support, fraud detection, autonomous system control, and environmental monitoring, while also guiding future research toward underutilized approaches such as reinforcement learning and explainable AI.

Abstract

Machine learning (ML) has become a central enabler of data-driven decision-making across smart manufacturing, healthcare, finance, computer vision, robotics, and environmental sustainability. Despite the rapid growth of ML applications, existing review studies remain largely domain-specific and provide limited cross-domain synthesis of methodological trends, deployment challenges, and emerging research directions. This systematic literature review aims to provide a comprehensive and comparative analysis of ML applications across seven high-impact domains while identifying dominant learning paradigms, implementation challenges, and future research opportunities. Following the PRISMA 2020 guidelines, peer-reviewed studies published between 2015 and 2025 were systematically collected from major scientific databases, including ScienceDirect, IEEE Xplore, SpringerLink, Wiley Online Library, MDPI, and Web of Science. Studies were screened using predefined inclusion and exclusion criteria and categorized according to application domain, ML paradigm, algorithm type, data characteristics, and deployment context. The findings indicate that supervised learning and deep learning dominate most application areas, with convolutional neural networks emerging as the primary approach for image-based and perception-driven tasks. Reinforcement learning, although highly promising for sequential decision-making and adaptive control, remains comparatively underutilized due to safety, computational, and deployment constraints. Across domains, recurring challenges include data quality, interpretability, scalability, model robustness, computational requirements, and ethical considerations. Overall, this review provides a structured cross-domain synthesis of ML applications and highlights the growing importance of explainable, trustworthy, and deployable AI systems for future intelligent and sustainable technologies.



Academic Editors: Douglas O'Shaughnessy and Roberto Zivieri

Received: 5 May 2026

Revised: 31 May 2026

Accepted: 27 June 2026

Published: 1 July 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: systematic literature review; machine learning applications; explainable Artificial Intelligence (XAI); cross-domain analysis; intelligent systems; deep learning; reinforcement learning

1. Introduction

1.1. Background and Rationale

Machine learning (ML) has emerged as a foundational technology for intelligent data-driven decision-making across a wide range of domains, including smart manufacturing, healthcare, finance, computer vision, robotics, and environmental sustainability. Advances in computational power, data acquisition technologies, cloud computing, and artificial intelligence (AI) algorithms have accelerated the adoption of ML systems capable of performing prediction, classification, optimization, anomaly detection, and autonomous decision-making tasks. Unlike traditional rule-based systems, ML models learn patterns and relationships directly from data, enabling adaptive and scalable solutions for increasingly complex real-world problems [1].

The proliferation of cyber-physical systems, IoT devices, and large-scale datasets has established ML as central to modern intelligent systems across manufacturing, healthcare, finance, computer vision, robotics, and environmental monitoring [2,3]. Despite this rapid expansion, existing reviews remain largely domain-specific and descriptive, lacking transparent cross-domain synthesis of methodological trends, deployment challenges, and paradigm comparisons [4]. Issues including explainability, scalability, reinforcement learning underutilization, and ethical deployment are addressed inconsistently across the literature. This SLR synthesizes peer-reviewed studies (2015–2025) across seven domains following PRISMA 2020 guidelines, providing a structured comparative assessment to support the development of robust, explainable, and deployable ML systems.

A critical gap in the existing literature is the failure to interrogate why certain paradigms persist across domains despite demonstrably superior alternatives, and what structural barriers, economic, regulatory, or epistemic, sustain these patterns. For instance, the continued dominance of supervised learning in healthcare and finance is not simply a matter of technical suitability; it reflects the legal and institutional constraints of regulated industries, where model auditability and liability are primary concerns that fundamentally constrain algorithm selection independent of predictive performance [5–8].

Existing domain-specific reviews seldom articulate these boundary conditions, thereby presenting a distorted picture of ML's applicability and overstating the transferability of methods across contexts. Furthermore, the comparative underutilization of reinforcement learning, observed consistently across six of the seven domains reviewed, represents not merely a research gap but a structural failure mode: RL's sample inefficiency and safety-critical exploration behavior create deployment barriers that no incremental algorithmic improvement alone can resolve, and which require coordinated advances in simulation fidelity, safe exploration theory, and domain-specific reward engineering [9].

1.2. Research Objectives and Research Questions

Specifically, this study addresses the following research questions:

- RQ1: Which machine learning paradigms and algorithms are most widely adopted across different application domains?
- RQ2: What are the major technical, operational, and deployment challenges associated with ML implementation across domains?

- RQ3: How are explainability, robustness, scalability, and trustworthiness addressed in existing ML applications?
- RQ4: What emerging trends and future research opportunities can be identified from the cross-domain synthesis of ML studies?

2. Materials and Methods

2.1. Review Protocol and Reporting Guidelines

This study was conducted as a systematic literature review (SLR) following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) guidelines to ensure methodological transparency, reproducibility, and structured reporting. The review process followed the four primary PRISMA phases: identification, screening, eligibility assessment, and inclusion. The PRISMA framework was adopted to systematically identify, evaluate, and synthesize peer-reviewed studies related to machine learning (ML) applications across multiple domains. The review protocol was designed to establish consistent procedures for literature searching, study selection, data extraction, and qualitative synthesis. The review was not formally registered in an external review registry.

The study selection process is illustrated in the PRISMA 2020 flow diagram below (Figure 1), summarizing the number of records identified, screened, assessed for eligibility, and ultimately included in the review.




Figure 1. PRISMA 2020 flow diagram—Study selection process.

2.2. Eligibility Criteria

Eligibility criteria were established prior to the literature search to ensure consistency and relevance in study selection. Studies were included if they satisfied the following criteria:

- Published in peer-reviewed journals or conference proceedings;
- Written in the English language;

- Published between 2015 and 2025;
- Focused on real-world applications of machine learning;
- Addressed at least one of the following domains: smart manufacturing, additive manufacturing, healthcare, finance and banking, computer vision, robotics and autonomous systems, or environmental sustainability;
- Investigated supervised learning, unsupervised learning, deep learning, reinforcement learning, or hybrid ML approaches;
- Provided sufficient methodological and application-related detail.

Studies were excluded if they:

- Were editorials, opinion papers, tutorials, white papers, theses, or non-peer-reviewed publications;
- Focused solely on theoretical algorithm development without application context;
- Lacked sufficient methodological clarity or application detail;
- Were duplicate publications;
- Fell outside the specified publication timeframe.

The literature search covered six databases, ScienceDirect, IEEE Xplore, SpringerLink, Wiley Online Library, MDPI, and Web of Science, using Boolean keyword combinations pairing ML paradigm terms with each domain application (e.g., ‘machine learning’ AND ‘predictive maintenance’; ‘deep learning’ AND ‘medical imaging’). Records were deduplicated, screened by title/abstract, then assessed by full text. A qualitative taxonomy-driven synthesis was adopted given the heterogeneity of datasets, metrics, and domains, focusing on dominant paradigms, deployment challenges, explainability considerations, and emerging research gaps.

2.3. Domain-Specific Applications

2.3.1. Application of ML in Smart Manufacturing, Additive Manufacturing, and Industry 4.0

Machine learning (ML) has become a central enabler of Industry 4.0 by supporting data-driven automation, monitoring, and optimization in increasingly complex manufacturing systems. The rapid growth of ML adoption in manufacturing is closely linked to the proliferation of big data generated by cyber-physical systems, smart sensors, and interconnected production platforms. While this data presents significant opportunities for enhanced decision-making, it also introduces challenges related to high dimensionality, heterogeneity, uncertainty, and system dynamics [2].

Manufacturing environments often suffer from incomplete, noisy, redundant, or irrelevant data, which can degrade ML model performance if not properly addressed [10]. In addition, the widespread use of cloud-based storage and wireless communication for remote monitoring and control raises critical concerns related to data privacy, cybersecurity, and system reliability. These challenges are further exacerbated when data are distributed across multiple physical locations, requiring robust communication infrastructures to support integrated learning and real-time decision-making [10].

Supervised learning, SVMs, decision trees, random forests, dominates predictive maintenance, quality prediction, and fault classification, driven by labeled sensor data availability [2,11]. Bertolini et al. [4] identified a sharp post-2015 increase in deep learning adoption, with CNNs emerging as the primary approach for visual inspection and defect detection. Reinforcement learning has been explored for adaptive scheduling and process optimization but remains limited by training costs and safety constraints [10,12]. For Additive Manufacturing (AM), supervised learning predicts surface roughness, dimensional accuracy, and defect occurrence from process parameters [13]. CNNs underpin layer-wise in-situ monitoring and melt-pool analysis, while the domain is characterized by small

datasets, limited benchmarks, and predominantly proof-of-concept deployments, with reinforcement learning offering underexplored potential for closed-loop process control [14,15].

Looking ahead, future trends emphasize the development of adaptive and autonomous ML algorithms capable of learning under uncertainty and dynamic conditions. Hybrid approaches that combine multiple learning paradigms are gaining traction in flexible manufacturing systems [16]. However, a persistent barrier to broader adoption remains the shortage of skilled professionals in data science, optimization, and industrial AI, which continues to limit the effective deployment of ML technologies in manufacturing environments [14].

The literature highlights several persistent and emerging challenges in modern manufacturing systems [14–19], including:

- noisy industrial sensor data;
- limited labeled datasets;
- model generalization across operating conditions;
- real-time deployment constraints.

These challenges collectively reflect the increasing complexity and dynamism of manufacturing systems, which are further intensified by uncertainty in global business environments. Complexity arises not only within production operations but also across product design lifecycles and inter-organizational business processes. Moreover, the widespread deployment of sensors, cyber-physical systems, and Industrial Internet of Things (IIoT) technologies has led to an unprecedented growth in manufacturing data. While this data richness creates significant opportunities, it also introduces major challenges related to high dimensionality, heterogeneity, noise, data quality, and the computational intractability of many manufacturing optimization problems [20].

2.3.2. Application of ML in the Finance & Banking

Empirical finance research has traditionally focused on modeling economic relationships and estimating the direction and magnitude of causal effects, such as the influence of firm characteristics on capital structure or the impact of regulatory changes on market behavior. Conventional econometric approaches prioritize inference and hypothesis testing, whereas machine learning (ML) emphasizes predictive performance and the discovery of complex patterns within large and high-dimensional financial datasets [21]. Rather than explicitly modeling causality, ML techniques learn functional relationships from data to generate predictions for new observations based on observed features.

Supervised methods such as logistic regression, SVMs, and ensemble methods (Random Forest, Boosted Regression Trees), dominate credit scoring, fraud detection, and risk assessment, collectively accounting for the majority of reviewed applications and favored for interpretability and regulatory compatibility [21]. Linear models retain a share where transparency is legally required. Reinforcement learning is more prominent in finance than in most other domains, applied to portfolio optimization and algorithmic trading via Deep Q-Networks and policy gradient methods, though deployment is constrained by market non-stationarity and risk management requirements [21,22]. DNNs lead deep learning applications for pattern recognition and market forecasting; RNNs appear least frequently despite their alignment with financial time series attributed to training complexity and hyperparameter sensitivity by Hoang and Wiegatz [21]. Persistent challenges include adversarial fraud evolution, model bias, and the need for fair, continually adaptive ML systems.

2.3.3. Application of ML in the Computer Vision

Machine learning (ML) is the foundational driver of modern computer vision systems, enabling automated interpretation of visual data across applications such as object detection, image classification, visual inspection, and scene understanding. Deep learning, particularly convolutional neural networks (CNNs), dominates this domain due to its ability to learn hierarchical spatial features directly from raw image data [23]. Supervised learning remains the most widely used paradigm, supported by the availability of large labeled image datasets, while unsupervised and self-supervised learning methods are gaining attention as viable alternatives to reduce reliance on costly manual labeling [24].

CNNs dominate across object detection, classification, and segmentation tasks, with Vision Transformers (ViT) increasingly competitive on large-scale benchmarks [23,25–27].

Key barriers include the labor-intensive nature of data labeling, the majority of real-world visual data is unstructured and unlabeled, high computational requirements for training, and concerns about synthetic dataset reliability [23,28,29]. Self-supervised approaches are gaining traction as scalable alternatives to manual annotation. Reinforcement learning has limited but emerging relevance in active perception and vision-based control. Explainability, adversarial robustness, and bias mitigation remain unresolved research priorities, while lightweight architectures (MobileNet, EfficientNet) address deployment constraints in resource-limited settings [25,26].

2.3.4. Application of ML in the Autonomous Systems & Robotics

The integration of artificial intelligence (AI), machine learning (ML), and deep learning (DL) has fundamentally transformed autonomous systems and robotics, enabling intelligent perception, control, and decision-making in complex and dynamic environments. These technologies allow robotic systems to automate tasks, adapt to changing conditions, and operate with increased precision, safety, and efficiency across industrial, service, and autonomous applications [30].

Supervised learning supports perception tasks, object recognition, localization, sensor fusion, using labeled visual and sensory data [31]. CNNs and DNNs dominate navigation, mapping, and motion planning, extracting spatial and temporal features from multimodal sensor inputs [31]. Reinforcement learning shows higher adoption in robotics than in most other reviewed domains, providing a principled framework for sequential control through environmental interaction [9]. However, RL deployment remains constrained by high sample complexity, simulation dependency, and the sim-to-real transfer gap [9,32]. Critical challenges include real-time computational demands, adaptability to dynamic environments, safety in human-robot interaction, and ethical implications of autonomous decision-making [33]. Future directions include safe RL, lifelong learning, and explainable autonomous control systems.

The use of AI and ML in robotics reduces labor dependency and maintenance expenditures, improving economic feasibility for industrial applications [34].

While AI, ML, and DL offer substantial benefits to robotics, they also present several significant challenges. One of the foremost issues is the requirement for large volumes of high-quality data to effectively train these algorithms. Data collection, labeling, and annotation are often costly and time-consuming, and the resulting datasets can be noisy or biased, factors that undermine model accuracy and reliability. This challenge is amplified in robotics, where acquiring consistent and relevant data is inherently difficult due to the complexity of real-world environments. Moreover, many robotics applications demand real-time data processing, which is computationally intensive and often requires specialized hardware [31]. The need for high processing power to handle massive data volumes, build predictive models, and make real-time decisions is at odds with the energy and computing

limitations of robotic platforms [35]. Another major hurdle is adaptability. Robots are expected to operate in dynamic, unpredictable environments, which necessitates the ability to learn from new experiences and adapt behavior accordingly. Designing AI/ML/DL models that can generalize beyond their training data and handle novel scenarios remains a considerable technical challenge. Safety is also a critical concern. As robots become more autonomous and increasingly interact with humans, ensuring safe and reliable operation becomes essential. AI systems must be able to detect and respond to hazards, avoid collisions, and function correctly in diverse environments [36,37]. Achieving this requires robust algorithms and high-fidelity sensors capable of accurately perceiving and interpreting the surroundings [38]. Beyond technical challenges, there are also ethical and societal implications. Concerns include the displacement of jobs due to automation, the potential for algorithmic bias, and the misuse of robotics in harmful applications such as military operations or surveillance [33]. These issues underscore the need for ethical frameworks and policy regulations to guide the responsible development and deployment of intelligent robotic systems.

2.3.5. Application of ML in the Healthcare

Machine learning (ML) has gained substantial traction in healthcare over the past decade, driven by advances in computational power, increased availability of large-scale medical data, and the digitization of healthcare systems. ML has contributed to automating clinical workflows, enhancing diagnostic accuracy, and supporting data-driven clinical decision-making. Its impact is particularly evident in medical imaging, electronic health record (EHR) analysis, disease risk prediction, and mobile health (mHealth) applications, where complex and high-dimensional data exceed the analytical capacity of traditional biostatistical methods.

Supervised learning represents the most widely adopted ML paradigm in healthcare, supporting tasks such as disease classification, outcome prediction, and clinical decision support. Models trained on structured and unstructured EHR data can identify disease-associated variables, uncover nonlinear relationships, and improve predictive accuracy for diagnosis and prognosis [5]. In medical imaging, deep learning, especially convolutional neural networks (CNNs), has achieved state-of-the-art performance in applications such as organ identification, tumor detection, lung disease classification, brain tumor segmentation, and image reconstruction, significantly advancing radiology and diagnostic imaging workflows [39].

Unsupervised learning techniques play a complementary role by enabling patient stratification, phenotyping, and exploratory analysis of complex biomedical datasets. These methods facilitate the discovery of latent disease subtypes and patterns that are not explicitly labeled, supporting personalized medicine and population health analysis. Reinforcement learning has also been explored for treatment optimization, adaptive therapy planning, and personalized intervention strategies. However, its real-world deployment remains limited due to ethical considerations, patient safety concerns, data scarcity, and the difficulty of validating sequential decision-making models in clinical settings.

2.3.6. Application of ML in the Environmental and Sustainability

Climate change, rapid urbanization, and population growth have intensified environmental pressures, posing significant challenges to the sustainability of urban and natural systems. In response, machine learning (ML) has increasingly been integrated with remote sensing and environmental monitoring technologies to support data-driven analysis, prediction, and decision-making. These approaches enable large-scale assessment of envi-

ronmental conditions, facilitate early warning systems, and support the development of mitigation and adaptation strategies.

A substantial body of research combines ML with remote sensing data such as satellite imagery, aerial photography, and LiDAR, to analyze land-use change, urban expansion, air quality, climate patterns, and ecosystem dynamics. These data sources are widely used due to their broad spatial coverage, high resolution, and cost-effectiveness. However, existing studies remain geographically concentrated in Asia, Europe, and North America, reflecting disparities in technological infrastructure, data availability, and governance frameworks. This concentration underscores the need for more inclusive and responsible ML deployment in data-scarce and underrepresented regions.

Supervised learning dominates environmental and sustainability applications, driven by the availability of labeled datasets for tasks such as land-cover classification, air quality prediction, and climate variable estimation [36]. Deep learning approaches, particularly convolutional neural networks, are extensively applied to satellite image analysis, object detection, and spatiotemporal modeling, enabling high-accuracy interpretation of complex environmental data. Unsupervised learning methods complement these efforts by supporting pattern discovery, clustering, and dimensionality reduction in large, heterogeneous environmental datasets, facilitating exploratory analysis and knowledge extraction [40].

Despite methodological advances, many ML-driven environmental systems continue to rely heavily on human oversight during critical stages, including image acquisition, data validation, and final decision-making. Moreover, existing research often prioritizes physical and spatial variables, such as land use, vegetation cover, and atmospheric composition, while underrepresenting socioeconomic factors. Integrating multi-source data that combines environmental, demographic, and economic information remains a key challenge for achieving holistic sustainability assessments.

The comparative analysis revealed both shared and domain-specific trends in machine learning adoption and deployment. Comparative Cross-Domain Synthesis is presented in Table 1.

Table 1. Comparative cross-domain synthesis.

Domain	Dominant Algorithms	Common Dataset Types	Major Challenges	Explainability Adoption	RL Adoption
Manufacturing	RF, SVM, ANN, CNN	Sensor data, process data	Noise, real-time constraints	Moderate	Low
Healthcare	CNN, RF, Logistic Regression	Medical images, EHRs	Privacy, regulation, imbalance	High	Low
Finance	XGBoost, DNN, Logistic Regression	Transactions, market data	Security, interpretability	High	Moderate
Computer Vision	CNN, Transformers	Images, video	Computational cost	Low–Moderate	Low
Robotics	RL, Deep RL, CNN	Sensor and simulation data	Safety, deployment complexity	Moderate	High
Environmental Systems	CNN, LSTM, RF	Climate and geospatial data	Data heterogeneity	Moderate	Low

Across domains, supervised learning remained the most mature and widely deployed paradigm, while deep learning dominated perception-based applications involving image and sequential data. Reinforcement learning demonstrated strong promise for adaptive decision-making but remained underutilized outside robotics and autonomous control systems.

Several common deployment challenges emerged across domains, including:

- Salability limitations;
- Lack of explainability;
- Computational cost;
- Data quality concerns;
- Robustness under changing operational conditions;

- Ethical and regulatory constraints.

The comparative synthesis further highlighted the growing importance of explainable, trustworthy, and deployable AI systems capable of balancing predictive performance with interpretability, robustness, and operational reliability.

2.4. Comparative Analysis Across Domains

This study's cross-domain analysis of Machine Learning (ML) applications reveals several key insights regarding the maturity, diversity, and research focus of ML integration across major industrial and societal sectors. The results demonstrate that while ML is broadly adopted, its depth, transparency, and methodological rigor vary substantially among domains. The pie chart in Figure 2, illustrates the distribution of ML Applications across various Domains and Table 2 represents the common ML Algorithms by Domain. Figure 3 presents the Heatmap of ML Paradigm Frequencies Across Domains.




Figure 2. Distribution of ML applications across various domains.

Table 2. Common machine learning algorithms by domain.

Domain	Frequently Used Machine Learning Algorithms
Industry 4.0 & Manufacturing	Random Forest, Support Vector Machines (SVM), Gradient Boosting, Convolutional Neural Networks (CNN), Long Short-Term Memory (LSTM)
Additive Manufacturing	Random Forest, Support Vector Machines (SVM), Gaussian Process Regression, CNN, LSTM
Finance & Banking	XGBoost, Random Forest, Logistic Regression, Neural Networks, Support Vector Machines (SVM)
Computer Vision	Convolutional Neural Networks (CNN), Vision Transformers (ViT), Autoencoders, U-Net, YOLO
Autonomous & Robotics	Reinforcement Learning (Q-learning, Deep Q-Networks), Policy Gradient Methods, CNN, LSTM
Healthcare	Logistic Regression, Random Forest, CNN, Recurrent Neural Networks (RNN), XGBoost
Environmental and Sustainability	Linear Regression, Random Forest, Gradient Boosting, LSTM, Gaussian Processes




Figure 3. Heatmap for frequency of ML paradigms across domains.

2.4.1. Applications of ML Methods Across Domains

Figures 4–10, represents the Pareto Chart of ML Methods in different Domains. According to figures, in the Environmental & Sustainability domain, supervised learning accounts for the largest share of reviewed ML applications, reflecting the widespread availability of labeled remote sensing, climate, and environmental monitoring datasets used in tasks such as air quality prediction, land-cover classification, and climate variable estimation [36]. This dominance is consistent with findings by Linaza et al. [36] and Garg [37], who identified supervised approaches as the primary paradigm in precision agriculture and environmental monitoring due to their strong performance on well-defined prediction targets. Deep learning constitutes the second largest share, with Convolutional Neural Networks (CNNs) and Long Short-Term Memory (LSTM) networks applied extensively to satellite image analysis, spatiotemporal forecasting, and object detection in geospatial data [25,26]. The complexity and high dimensionality of remote sensing data, including multispectral imagery, LiDAR point clouds, and temporal climate records, make deep architectures particularly well-suited to this domain. Reinforcement learning (RL) remains substantially underrepresented in environmental applications, accounting for a small fraction of reviewed studies. However, its theoretical potential for adaptive energy management, dynamic resource allocation, and smart grid optimization is well-established [36]. The limited adoption of RL is attributable to challenges in defining appropriate reward functions for continuous environmental systems, the high computational cost of training, and the difficulty of validating RL policies in safety-critical environmental contexts. Unsupervised learning plays a complementary rather than primary role, typically applied to clustering of land-use patterns, dimensionality reduction in high-dimensional geospatial datasets, and exploratory analysis where labeled training data are scarce.




Figure 4. Pareto chart of ML methods in environmental & sustainability.




Figure 5. Pareto chart of ML methods in Industry 4.0 & manufacturing.




Figure 6. Pareto chart of ML methods in finance.




Figure 7. Pareto chart of ML methods in autonomous systems and robotics.




Figure 8. Pareto chart of ML methods in additive manufacturing.




Figure 9. Pareto chart of ML methods in healthcare.




Figure 10. Pareto chart of ML methods in computer vision.

In the Finance and Banking domain, reinforcement learning accounts for a notably higher proportion of reviewed studies relative to most other domains, driven by its suitability for sequential trading decisions, portfolio optimization, and adaptive strategy development under dynamic market conditions [21,41]. Deep reinforcement learning variants, including Deep Q-Networks and policy gradient methods, are applied extensively to algorithmic trading and quantitative finance problems, where the capacity to model temporal reward structures provides a clear advantage over purely supervised formulations [21]. Supervised learning retains a significant presence, particularly for credit risk prediction, fraud detection, and return forecasting tasks, where labeled historical datasets enable reliable training of classification and regression models. Deep learning architectures, including Deep Neural Networks (DNNs) and CNNs, have gained traction for processing unstructured textual data such as earnings announcements and sentiment signals from financial news, though their adoption remains less widespread than in perception-driven domains. The relatively balanced distribution of ML paradigms in finance, contrasting with the supervised-learning dominance observed in manufacturing and healthcare, reflects the domain's unique combination of structured transaction data, sequential decision problems, and high tolerance for algorithmic complexity.

In the Autonomous Systems and Robotics domain, deep learning constitutes the dominant ML paradigm, accounting for the largest share of reviewed applications. This predominance reflects the central role of perception tasks, including visual object detection, semantic segmentation, and sensor fusion, in autonomous systems, for which CNNs and DNNs are architecturally well-suited [31]. Convolutional architectures in particular underpin the majority of vision-based control and navigation applications reviewed, consistent with their proven performance on image recognition benchmarks and their capacity to learn robust spatial feature representations directly from raw sensor data [31]. Reinforcement learning maintains a meaningful but comparatively smaller share, reflecting its application to sequential decision-making and adaptive robot control problems, including manipulation, locomotion, and autonomous navigation [9]. Its adoption remains constrained by well-documented practical barriers: high sample complexity, the need for simulation environments to generate sufficient training experience, and significant challenges in transferring learned policies to real-world deployment settings [9]. Supervised learning represents the smallest paradigm share in this domain, reflecting the limited applicability of purely label-driven prediction approaches to the continuous-action, high-dimensional control problems that characterize robotics.

In the Additive Manufacturing domain, supervised learning leads the distribution of reviewed ML applications, followed by deep learning, with unsupervised and reinforcement learning accounting for only marginal shares. The dominance of supervised methods reflects the structured nature of many AM prediction tasks, such as dimensional accuracy prediction, process parameter optimization, and defect classification, which pair well with labeled training datasets derived from experimental print runs [13]. However, the availability of sufficiently large and diverse labeled datasets remains a recognized challenge in this domain, particularly for novel materials and non-standard build geometries [14]. Deep learning methods, primarily CNNs and ANNs, are applied to in-situ monitoring tasks including layer-wise defect detection and thermal process characterization, where the capacity for automated feature extraction from image and sensor data provides a measurable advantage over hand-engineered feature approaches [13]. The modest but growing presence of unsupervised methods, primarily PCA and K-Means Clustering, reflects their role in print anomaly detection and exploratory analysis of process parameter spaces where labeled data are unavailable.

In the Healthcare domain, supervised learning is the overwhelmingly dominant paradigm, consistent with the availability of large, well-curated clinical datasets, including diagnostic labels, treatment outcomes, and annotated medical images, that enable reliable model training for classification and prediction tasks [5]. Logistic regression, random forests, and support vector machines are extensively applied to structured electronic health record (EHR) data for disease prediction, risk stratification, and clinical decision support, where their interpretability supports regulatory compliance and clinician trust [39]. Deep learning, particularly CNNs, accounts for the second largest share, driven primarily by medical imaging applications in radiology and pathology, where convolutional architectures demonstrate strong performance on tasks such as tumor detection, lesion segmentation, and diagnostic classification [5]. The moderate but meaningful adoption of deep learning in healthcare reflects both its predictive power and the growing availability of large annotated imaging datasets such as those from MIMIC and similar repositories. Recurrent neural networks and other sequence models remain comparatively underrepresented despite their theoretical suitability for temporal EHR data and longitudinal patient monitoring, suggesting a continuing research gap in sequential clinical modelling [7,8]. Unsupervised methods are applied selectively to patient clustering, dimensionality reduction in high-dimensional genomic data, and anomaly detection in monitoring streams, though their deployment in clinical settings is constrained by the interpretability requirements of healthcare applications.

In the Computer Vision domain, deep learning accounts for the overwhelming majority of reviewed applications, to a degree that distinguishes this domain from all others in the synthesis. This near-exclusive reliance on deep architectures reflects the fundamental nature of visual data processing: the hierarchical, multi-layer feature representations learned by CNNs are architecturally well-matched to the spatial structure of image data, enabling end-to-end learning from raw pixels without manual feature engineering [23]. CNNs dominate across object detection, image classification, and segmentation tasks, while transformer-based architectures, including Vision Transformers (ViT), are increasingly adopted for tasks requiring long-range spatial context modelling [25,26]. The limited presence of supervised non-deep methods, including SVMs and logistic regression, is confined to lower-complexity classification settings where smaller datasets or computational constraints preclude the use of deep architectures. Reinforcement learning has minimal representation in the reviewed computer vision literature, with its application largely confined to active perception and attention-driven image analysis tasks. The dominance of deep learning in this domain also reflects an important limitation: the computational and data requirements of deep vision

models create significant barriers for deployment in resource-constrained environments, a challenge that emerging lightweight architectures such as MobileNet and EfficientNet are beginning to address [25,26].

2.4.2. Applications of ML Algorithms by Method Type Across Domains

Figures 11–17, represents the Stacked Bar Chart of ML Algorithms by Method Type in different Domains. According to figures, in the Environmental and Sustainability domain, Support Vector Machines (SVMs) and Random Forests lead among supervised algorithms, reflecting their strong track record in land-cover classification, air quality prediction, and environmental risk modelling tasks where labeled remote sensing datasets are available [36]. Their prevalence over simpler methods such as linear regression reflects the nonlinear and high-dimensional character of environmental feature spaces derived from satellite imagery and atmospheric measurement systems. Among unsupervised approaches, PCA and Gaussian Mixture Models (GMMs) are the primary methods applied, principally for dimensionality reduction of hyperspectral and multispectral datasets and for clustering spatiotemporal environmental monitoring records. Within deep learning, CNNs lead the distribution, consistent with their strong performance on spatial image analysis tasks such as land-use change detection and urban expansion mapping [25,26]. LSTM networks and RNNs account for a secondary deep learning share, applied to temporal forecasting of climate variables, energy demand, and atmospheric pollution concentrations where sequential dependencies in time-series data are critical. Reinforcement learning shows a limited but notable presence in manufacturing and Industry 4.0 applications within this grouping, applied primarily to dynamic production scheduling, inventory management, and process optimization under variable demand conditions [42]. The integration of ML with IoT sensor networks and digital twin frameworks is an emerging trend in this domain, with several reviewed studies demonstrating that combining real-time sensor streams with predictive models substantially improves fault detection lead times and reduces unplanned downtime [11].

In the Finance and Banking domain, the distribution of supervised algorithms reveals a concentration around high-performing nonlinear and ensemble methods. Support Vector Machines (SVMs), Random Forests (RF), and Boosted Regression Trees (BRT) collectively account for a substantial majority of the supervised learning applications reviewed, reflecting their robust generalization performance on financial prediction tasks such as credit default forecasting, return prediction, and fraud classification, where complex nonlinear relationships between features are common [21]. Traditional linear models, including logistic regression and linear regression, account for a smaller but non-negligible share, typically applied in settings where interpretability and regulatory transparency are paramount, such as credit scoring subject to fair lending requirements [21]. Among unsupervised methods, K-Means Clustering and Gaussian Mixture Models (GMMs) are the primary algorithms identified, applied to customer segmentation and portfolio clustering tasks; their limited overall frequency in the reviewed literature indicates that unsupervised learning is not yet widely deployed in mainstream financial ML workflows. Within deep learning, Deep Neural Networks (DNNs) account for the largest share, followed by CNNs and ANNs applied to text-based sentiment analysis and multi-modal financial data fusion tasks. Recurrent Neural Networks (RNNs) appear least frequently despite their theoretical alignment with sequential financial time series, a gap noted also by Hoang and Wiegatz [22], who observed that the complexity of RNN training and hyperparameter sensitivity may limit their adoption in practitioner-oriented financial research. The overall algorithm distribution in finance reflects the domain's dual emphasis on predictive accuracy and model interpretability, with ensemble methods striking a practical balance between the two objectives.




Figure 11. Stacked bar chart of ML algorithms by method type in environmental & sustainability.




Figure 12. Stacked bar chart of ML algorithms by method type in Industry 4.0 & manufacturing.




Figure 13. Stacked bar chart of ML algorithms by method type in finance.




Figure 14. Stacked bar chart of ML algorithms by method type in autonomous systems and robotics.




Figure 15. Stacked bar chart of ML algorithms by method type in additive manufacturing.




Figure 16. Stacked bar chart of ML algorithms by method type in healthcare.




Figure 17. Stacked bar chart of ML algorithms by method type in computer vision.

In the Autonomous Systems and Robotics domain, the algorithm distribution is strongly concentrated in deep learning, with CNNs and DNNs together accounting for the large majority of reviewed applications. CNNs are deployed primarily for perception tasks, including object detection, obstacle recognition, and scene understanding, which are foundational to autonomous navigation and inspection systems [31]. DNNs serve broader function-approximation roles in control pipelines where sensor inputs must be mapped to continuous action spaces. Supervised learning is confined almost exclusively to SVM-based applications in this domain, reflecting its limited applicability to the continuous, high-dimensional control problems characteristic of robotics. Reinforcement learning, represented primarily by Q-Learning and Deep Q-Learning methods, accounts for a smaller but analytically significant share, applied to manipulation, locomotion, and adaptive control problems [9]. The underrepresentation of RL relative to deep learning is consistent with the substantial practical barriers to its deployment in physical systems, including sample inefficiency, safety constraints during exploration, and the sim-to-real transfer gap that remains an open research challenge [9]. Notably, the reviewed literature shows minimal use of Artificial Neural Networks (ANNs) in this domain, a pattern that distinguishes robotics from manufacturing and healthcare, where ANNs remain prevalent for regression and classification tasks on structured sensor data.

In the Additive Manufacturing domain, deep learning architectures dominate the reviewed literature, with Convolutional Neural Networks (CNNs) applied primarily to in-situ defect detection and layer-wise print quality monitoring, and Artificial Neural Networks (ANNs) employed for process parameter prediction and performance modelling [13]. Among unsupervised methods, usage remains limited; where applied, Principal Component Analysis (PCA) and K-Means Clustering are the predominant choices, typically serving dimensionality reduction and print anomaly clustering functions [14]. Overall, ML adoption in additive manufacturing remains at a comparatively early stage of maturity relative to domains such as manufacturing and healthcare, a pattern also observed by Jiang et al. [13], who noted that the field is characterized by small experimental datasets, limited benchmark standards, and a predominance of proof-of-concept studies over deployed systems. The predominance of supervised learning approaches, which require substantial volumes of labeled defect or process data, may constrain methodological diversity and hinder the exploration of more adaptive paradigms capable of operating in low-label or continually changing production environments [13]. Emerging research directions include reinforcement learning for self-optimizing print control strategies and unsupervised in-

process quality assessment frameworks capable of detecting novel defect types without prior annotation.

In the Healthcare domain, logistic regression and decision trees are the most frequently applied supervised algorithms, reflecting the field's strong emphasis on model interpretability and clinical transparency [5]. These methods offer predictive performance competitive with more complex approaches on structured tabular clinical data, while supporting explainable outputs that are essential for regulatory approval and clinical adoption [7,8]. Random forests and support vector machines also appear frequently, particularly for multi-class disease classification and risk stratification tasks using EHR-derived features. Among unsupervised approaches, K-Means Clustering and PCA are applied selectively, primarily for patient cohort discovery in population health studies and for feature reduction in high-dimensional genomic or proteomic datasets, though their deployment in direct clinical decision support remains limited due to interpretability concerns [39]. Within deep learning, CNNs account for the dominant share, concentrated in medical imaging applications including radiological diagnosis, histopathological image analysis, and retinal fundus assessment, where their spatial feature extraction capabilities significantly outperform traditional handcrafted feature approaches [5]. Recurrent Neural Networks and LSTM models are underrepresented relative to their theoretical suitability for longitudinal EHR and wearable sensor data, suggesting an important research opportunity for sequence-aware clinical modelling. The overall algorithm distribution reflects the particular demands of healthcare ML: predictive accuracy must be balanced against interpretability, regulatory compliance, and safety, a constraint that continues to favor traditional statistical and ensemble methods alongside selectively applied deep learning systems [7,8].

In the Computer Vision domain, CNNs and Deep Neural Networks (DNNs) collectively dominate the algorithm distribution to a degree not observed in any other domain in this review. CNNs are the primary architecture across object detection, image classification, semantic segmentation, and visual quality inspection tasks, underpinned by a large body of evidence demonstrating their superiority over hand-crafted feature methods on standard computer vision benchmarks [23]. Within the supervised non-deep category, Support Vector Machines and logistic regression retain a presence in lower-complexity recognition tasks and in settings where the limited size of training datasets makes deep architecture training impractical. The emergence of Vision Transformers (ViT) and attention-based architectures represents a significant methodological development within the period reviewed, with transformer models increasingly competitive with or superior to CNNs on large-scale image recognition benchmarks [25,26]. This trend suggests that the algorithmic landscape of computer vision is undergoing a paradigm shift, with implications for the types of pre-training, fine-tuning, and transfer learning strategies that will define the field in coming years [43].

2.4.3. Comparative Frequency of Machine Learning Paradigms Across Seven Application Domains

Figure 18 summarizes the distribution of supervised learning, unsupervised learning, deep learning, and reinforcement learning methods identified through the systematic literature review. The results highlight the dominance of supervised learning across all domains and the concentration of deep learning in data- and image-intensive domains such as healthcare, computer vision, and environmental sustainability, while reinforcement learning remains underutilized across most application areas.




Figure 18. Comparative frequency of ML paradigms across domains.

Figures 19–21 collectively confirm that algorithm selection is strongly domain-dependent. CNN usage peaks in Computer Vision and remains moderate in Environmental and Manufacturing domains; ANN/DNN models show consistent cross-domain applicability; SVMs and logistic regression retain higher shares in Healthcare and Finance where interpretability is prioritized; RNN adoption is uniformly low; and RL has the lowest overall representation, with meaningful presence only in Robotics and Finance.




Figure 19. Percentage breakdown of ML paradigms within each application domain.




Figure 20. ML algorithm frequency distribution per domain.




Figure 21. Radar chart of relative algorithm emphasis across domains.

3. Discussion

3.1. Principal Findings

This systematic literature review provides a comprehensive cross-domain synthesis of machine learning (ML) applications across seven major domains, including Industry 4.0 and manufacturing, additive manufacturing, healthcare, finance and banking, computer vision, robotics and autonomous systems, and environmental sustainability. The findings demonstrate that while ML adoption has expanded substantially across all domains, the depth of implementation, methodological maturity, and deployment readiness vary consid-

erably depending on application requirements, data availability, interpretability needs, and operational constraints.

Supervised learning and deep learning dominate across all domains due to labeled data availability and proven performance. SVMs, random forests, logistic regression, and ANN/DNN architectures provide strong results with manageable implementation complexity. CNNs lead in image-intensive domains, computer vision, healthcare imaging, environmental remote sensing, and autonomous perception, significantly outperforming traditional approaches on recognition and detection tasks [2,5,23]. RNNs remain underutilized relative to their suitability for temporal data across multiple domains. Reinforcement learning, despite strong theoretical potential for adaptive control and sequential decision-making, is substantially underutilized in practice, concentrated in robotics, autonomous systems, and finance, constrained by computational cost, training instability, simulation dependency, and safety requirements [9]. Environmental and sustainability applications show the strongest overall ML adoption. Healthcare and finance prioritize interpretability and regulatory compliance, often favoring explainable models despite the performance advantages of deep learning [5,7,8,21]. However, stating that supervised learning “dominates due to labeled data availability” without unpacking the implications obscures a deeper structural tension: the reviewed literature reveals that labeled data availability is itself a product of domain-specific historical investments in annotation infrastructure, not a natural property of the domain. In manufacturing, labeled sensor data is abundant precisely because predictive maintenance programs have systematically generated failure event logs over decades [2,11]; in additive manufacturing, the inverse holds, labeled defect data is scarce because the field lacks standardized benchmark datasets and industrial actors treat process data as proprietary [13,14]. This distinction has direct methodological consequences: methods well-suited to data-rich manufacturing environments cannot be straightforwardly transferred to additive manufacturing contexts without addressing the data generation pipeline itself. Similarly, the observation that healthcare and finance “prioritize interpretability” requires deeper unpacking. In healthcare, interpretability requirements are not merely a preference but are embedded in regulatory frameworks such as the EU’s Medical Device Regulation (MDR) and the FDA’s guidance on AI/ML-based Software as a Medical Device (SaMD), which impose explicit documentation and transparency obligations that constrain which model classes are deployable regardless of predictive accuracy [7,8]. In finance, interpretability mandates arise from fair lending legislation requiring institutions to provide adverse action notices to applicants, which necessitates models with auditable feature attribution, a requirement that deep neural networks cannot satisfy without post-hoc explanation methods whose reliability remains contested [21]. The cross-domain pattern of RL underutilization is similarly more complex than it appears: while computational cost is a real barrier, the more fundamental obstacle is the absence of validated simulation environments that faithfully model domain dynamics at sufficient fidelity for learned policies to transfer reliably. In robotics, this sim-to-real gap has received substantial research attention [9]; in healthcare and environmental applications, the development of such environments remains an open and largely unaddressed research problem.

3.2. Cross-Domain Methodological Patterns

Several important methodological patterns emerge across the reviewed literature. First, the majority of studies rely heavily on supervised learning approaches, reflecting the dependence of current ML systems on labeled datasets. While supervised methods provide reliable predictive performance, this reliance may constrain scalability and limit the discovery of hidden structures in unlabeled or data-scarce environments.

The pervasiveness of supervised learning across domains is not simply a reflection of its technical strengths but also an artifact of publication incentives and benchmark culture: studies reporting high accuracy on standard labeled benchmarks are more readily accepted, which systematically biases the literature toward supervised approaches and creates a self-reinforcing cycle that underrepresents the true potential of semi-supervised, self-supervised, and unsupervised methods [23,25,26]. This publication bias has measurable consequences: domains such as additive manufacturing, where labeled data is genuinely scarce, exhibit a mismatch between the dominant methodological paradigm (supervised learning) and the actual data conditions encountered in practice (small, unlabeled, or weakly annotated datasets) [13,14]. Addressing this disconnect requires not only methodological innovation but also deliberate changes in how systematic reviews and empirical studies are conducted, specifically by expanding eligibility criteria and reporting standards to include non-standard paradigms, and by emphasizing deployment realism over benchmark performance.

Second, deep learning adoption is highly domain-dependent. Domains characterized by image, video, or spatial data strongly favor CNN-based architectures, whereas structured-data domains such as finance and healthcare continue to utilize interpretable methods such as logistic regression, decision trees, and SVMs alongside deep learning approaches.

The domain-dependence of deep learning adoption reflects a deeper architectural alignment principle that the reviewed literature rarely makes explicit: CNNs exploit the translational invariance and spatial locality of image data through hierarchical convolution, which is precisely why they outperform dense architectures on vision tasks but provide no analogous inductive bias advantage on tabular or time-series data [23]. This architectural specificity has significant practical implications. The widespread adoption of CNNs in environmental remote sensing is well-founded, as satellite imagery exhibits exactly the spatial properties for which convolutional operators are designed [25,26]. However, their extension to non-image environmental data, such as meteorological station measurements or discrete air quality indices, is methodologically questionable, yet such extensions appear in the reviewed literature without sufficient justification. A more rigorous cross-domain framework would require researchers to articulate the inductive bias alignment between their chosen architecture and the structure of their data, rather than defaulting to deep learning as a performance-maximizing heuristic. The emergence of Vision Transformers (ViT) in computer vision [25,26] adds further analytical complexity: unlike CNNs, ViT models do not impose spatial locality constraints and instead learn long-range dependencies through self-attention, suggesting they may be better suited to domains with globally structured patterns, an implication that has not yet been systematically explored in healthcare imaging or geospatial analysis.

Third, explainability and interpretability considerations remain inconsistently addressed across domains. Although several studies incorporated explainable AI techniques such as SHAP and LIME, many high-performing deep learning systems continue to operate as black-box models with limited transparency. This issue is especially critical in high-stakes applications such as healthcare diagnosis, financial risk assessment, and autonomous systems, where accountability and trust are essential [44,45].

The inconsistency in XAI adoption identified across this review reflects not only a technical deficit but a conceptual one: the reviewed literature frequently conflates post-hoc explainability, techniques such as SHAP and LIME that generate approximate local explanations after prediction, with true model interpretability, which refers to the capacity of a model's internal structure to be directly understood by a domain expert [7,8]. This distinction is critical because post-hoc methods are approximations whose fidelity to the

underlying model's decision logic cannot be guaranteed, a limitation that is especially consequential in safety-critical domains. For instance, a SHAP explanation of a neural network's diagnostic prediction may accurately indicate which input features most influenced the output while nonetheless misrepresenting the actual computational pathway, creating a false sense of transparency that could mislead clinicians or regulators [7,8].

Furthermore, the reviewed literature shows that XAI adoption is primarily concentrated in domains where external pressures, regulatory mandates, clinical governance requirements, or litigation exposure, create explicit incentives for transparency, rather than reflecting a principled methodological commitment to interpretability. In computer vision and robotics, where such pressures are weaker, XAI adoption is lowest despite the high-stakes nature of autonomous decision-making in these domains. This pattern suggests that interpretability in ML applications is currently driven by compliance rather than by a mature understanding of its epistemic and ethical value, a gap that future research agendas must address directly.

Another important methodological observation is the relatively stable presence of ANN/DNN architectures across all domains, highlighting their versatility as general-purpose learning models. In contrast, RNN-based methods remain comparatively underrepresented, suggesting that temporal and sequential modeling continues to represent an underexplored research opportunity in several domains.

The underrepresentation of RNNs and LSTMs across the reviewed domains is particularly striking given the prevalence of temporal data in nearly all application areas, equipment sensor streams in manufacturing, electronic health records with longitudinal follow-up in healthcare, financial time series, climate monitoring records, and sequential robot trajectories in autonomous systems. The explanation for this gap is [46].

First, training RNNs on long sequences is computationally demanding and prone to vanishing gradient pathologies, whereas CNNs applied to temporal data through 1D convolutions provide a simpler and often comparably effective alternative [31].

Second, the literature reveals a pattern of methodological inertia: domains that developed their ML pipelines primarily around tabular or image data in the early 2010s have not systematically revisited those architectures as better-suited sequential models became available.

Third, the rapid rise of transformer-based architectures, originally developed for natural language processing but increasingly applied to time-series data through models such as the Temporal Fusion Transformer, is only beginning to appear in the reviewed literature, suggesting that the field is on the cusp of a paradigm shift in sequential modeling that this synthesis identifies as a high-priority future research direction [25,26].

The implications are domain-specific: in healthcare, the underuse of sequence models for longitudinal EHR data represents a missed opportunity for more accurate prognosis and treatment response prediction [7,8]; in environmental sustainability, LSTM-based spatiotemporal forecasting has demonstrated superior performance on climate variable prediction tasks yet remains a minority approach in the reviewed literature [36].

3.3. Limitations of Included Studies

Several methodological weaknesses and limitations were consistently identified across the reviewed literature.

Five recurring limitations were identified across the reviewed literature. (1) Absence of standardized evaluation metrics, most studies omitting F1-score, ROC-AUC, or benchmarking results, substantially limits cross-study comparability and reproducibility. (2) Reliance on small, imbalanced, or narrowly scoped datasets, particularly in healthcare, robotics, and environmental studies, reduces generalizability and is compounded by the lack of

external validation. (3) Strong dependence on supervised learning restricts exploration of unsupervised, self-supervised, and reinforcement learning methods better suited to dynamic or data-scarce environments. (4) Many studies prioritize predictive performance over explainability, transparency, and uncertainty quantification, critical in healthcare, finance, and autonomous systems where regulatory compliance and stakeholder trust are essential [5,7,8]. (5) ML research remains geographically concentrated in North America, Europe, and parts of Asia, leaving developing economies underrepresented and limiting global applicability.

These limitations are not independent; they form an interconnected cluster of methodological deficiencies with compounding effects on the validity and applicability of the reviewed findings. The absence of standardized evaluation metrics (limitation 1) is most damaging in conjunction with small, imbalanced datasets (limitation 2): when a study reports only accuracy on a dataset with 90% negative class prevalence, the apparent high performance may reflect the trivial classifier strategy of predicting the majority class, not genuine model learning, a well-documented pathology in medical ML research [5,21,27,47–49].

Without F1-score or ROC-AUC reporting, this cannot be detected by the reader, and such studies enter the corpus of reviewed evidence with inflated apparent performance. This interaction effect explains why the literature on healthcare ML in particular presents a more optimistic picture of model capability than clinical deployment outcomes would support. The geographic concentration of ML research (limitation 5) warrants more pointed analysis than is typically offered. The dominance of studies from North America, Europe, and East Asia reflects not only infrastructure advantages but also the dominance of these regions' institutions in high-impact journals and benchmark dataset creation, a feedback loop that systematically marginalizes findings from data-scarce regions where the need for deployable ML solutions may be most acute. In environmental sustainability, for example, the reviewed literature disproportionately addresses land-use and climate monitoring in temperate, data-rich geographies, while leaving semi-arid, tropical, and small-island developing state (SIDS) environments, which face the most severe climate change impacts, underrepresented [36]. Addressing this imbalance requires not only expanding research geographies but also developing ML frameworks robust to data scarcity and environmental conditions not well-represented in existing training corpora.

3.4. Implications for Practice and Future Research

The findings of this review highlight several important implications for both researchers and practitioners.

Future research should systematically integrate XAI techniques, SHAP, LIME, feature attribution, and intrinsically interpretable models, into ML workflows for healthcare, finance, and autonomous systems where accountability is non-negotiable. Safe reinforcement learning, offline RL, and uncertainty-aware learning represent critical development directions to bridge the gap between RL's theoretical potential and practical deployment. Interdisciplinary integration of ML with IoT, digital twins, edge computing, and physics-informed modeling can improve scalability and deployment reliability across domains [5,6,42]. Standardized reporting of datasets, preprocessing, hyperparameter settings, and evaluation metrics must become normative to enable reproducibility. Expanding ML research into underrepresented regions is essential for equitable technological development [5,6].

Table 3 provides a structured cross-domain overview of the key technical and operational challenges currently limiting ML deployment in each application domain, together with prioritized future research directions aimed at addressing those barriers and advancing toward more robust, explainable, and deployable ML systems.

Table 3. ML Research Challenges and Future Directions across Application Domains.

Domain	Key Challenges	Future Directions
Industry 4.0 & Manufacturing	Data heterogeneity, limited labels, model transferability	Physics-informed ML, federated learning, digital twin integration
Additive Manufacturing	Sparse data, process variability, real-time control	Reinforcement learning, hybrid ML-physics models
Healthcare	Interpretability, data privacy, regulatory constraints	Explainable AI, uncertainty-aware models, privacy-preserving ML
Finance	Bias, transparency, evolving fraud patterns	Fair and interpretable ML, continual learning
Computer Vision	Annotation cost, domain shift	Self-supervised learning, foundation models
Autonomous & Robotics	Safety, sample inefficiency, real-world deployment	Safe RL, sim-to-real transfer, lifelong learning
Environmental & Sustainability	Data sparsity, climate uncertainty	Transfer learning, probabilistic ML, multi-source data fusion

4. Conclusions

This systematic literature review synthesized ML applications across seven high-impact domains following PRISMA 2020 guidelines. Supervised learning and deep learning dominate due to labeled data availability and proven performance; CNNs lead image-intensive domains while ensemble methods prevail in structured-data settings. Reinforcement learning remains substantially underutilized despite strong potential for adaptive control and sequential decision-making, constrained by computational cost, safety concerns, and simulation dependency. Explainability and trustworthiness are growing priorities, particularly in healthcare and finance, yet XAI integration in deployed systems remains inconsistent. Recurring methodological gaps include insufficient performance reporting, limited external validation, and dataset scope constraints. This review contributes a structured cross-domain taxonomy and domain-specific research agenda, underscoring that future progress requires not only predictive accuracy improvements but transparent, equitable, and deployable ML systems capable of reliable real-world impact.

Three overarching conclusions emerge from this synthesis that extend beyond the enumeration of domain-specific findings. First, the cross-domain evidence reveals a systematic decoupling between methodological maturity and deployment maturity: domains such as computer vision exhibit high methodological sophistication, with transformer-based architectures, self-supervised pretraining, and benchmark-driven evaluation protocols, yet face significant deployment barriers due to computational resource requirements and the difficulty of generalizing models trained on curated benchmarks to noisy real-world distributions [23,25,26].

Conversely, manufacturing and environmental domains show comparatively lower methodological diversity but higher deployment prevalence, reflecting the pragmatic preference of industry practitioners for interpretable, reliable models over high-performing experimental ones [2,11].

This decoupling suggests that methodological progress in ML research and operational progress in ML deployment are not well-coupled, and that bridging this gap requires deliberate investment in deployment-oriented research, including studies of model degradation over time, retraining protocols, and integration with domain-specific operational systems, which the current literature addresses only superficially. Second, the review identifies a coherent and underappreciated pattern in the relationship between domain risk profiles and XAI adoption: high-stakes domains with explicit regulatory accountability frameworks (healthcare, finance) show higher XAI adoption, but this adoption is driven

by compliance pressure rather than by a principled understanding of what constitutes adequate explainability for a given decision context [5,7,8,50].

This compliance-driven adoption produces a troubling dynamic in which XAI techniques are applied as documentation artifacts rather than as tools for genuine understanding, a distinction that regulators and practitioners must actively engage with as AI governance frameworks mature. Third, the reinforcement learning gap identified across six of seven domains is not simply a research opportunity; it represents a fundamental bottleneck in the development of autonomous, adaptive ML systems capable of operating in dynamic real-world environments. Resolving this bottleneck requires coordinated advances across three dimensions: simulation environment development, safe exploration algorithms, and domain-specific reward formalization, none of which can be addressed through purely algorithmic research alone [9].

A comprehensive list of additional papers in each domain, with Authors, Publication Year, ML Method & Algorithms used, and Real Implication of Method & Algorithms in each Domain are provided as Supplementary Materials of this work. Future systematic reviews should track progress across these dimensions explicitly to assess whether the field is converging on practical RL deployment or whether the gap is widening as other paradigms continue to advance more rapidly.

Supplementary Materials: The following supporting information can be downloaded at: <https://www.mdpi.com/article/10.3390/app16136574/s1>.

Author Contributions: Conceptualization, N.S.; methodology, N.S.; software, N.S.; validation, N.S., S.E.K. and B.R.R.D.; formal analysis, N.S.; investigation, N.S.; resources, N.S.; data curation, N.S.; writing, original draft preparation, N.S.; writing, review and editing, N.S., S.E.K. and B.R.R.D.; visualization, N.S.; supervision, N.S.; project administration, N.S.; funding acquisition, N.S. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: Data sharing is not applicable to this article, as this is a review paper. All the references used for this paper is available on the Supplementary File S1.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Lecun, Y.; Bengio, Y.; Hinton, G. Deep learning. *Nature* **2015**, *521*, 436–444. [CrossRef] [PubMed]
2. Wuest, T.; Weimer, D.; Irgens, C.; Thoben, K.-D. Machine learning in manufacturing: Advantages, challenges, and applications. *Prod. Manuf. Res.* **2016**, *4*, 23–45. [CrossRef]
3. Jordan, M.I.; Mitchell, T.M. Machine learning: Trends, perspectives, and prospects. *Science* **2015**, *349*, 255–260. [CrossRef] [PubMed]
4. Bertolini, M.; Mezzogori, D.; Neroni, M.; Zammori, F. Machine Learning for industrial applications: A comprehensive literature review. *Expert Syst. Appl.* **2021**, *175*, 114820. [CrossRef]
5. An, Q.; Rahman, S.; Zhou, J.; Kang, J.J. A Comprehensive Review on Machine Learning in Healthcare Industry: Classification, Restrictions, Opportunities and Challenges. *Sensors* **2023**, *23*, 4178. [CrossRef] [PubMed]
6. Pan, L.; Li, G.; Zhu, T.; Liu, D.; Wang, Y.; Lu, Y. Physics-Informed Machine Learning in Design and Manufacturing: Status and Challenges. *J. Comput. Inf. Sci. Eng.* **2025**, *25*, 120804. [CrossRef]
7. Abdar, M.; Khosravi, A.; Islam, S.M.S.; Acharya, U.R.; Vasilakos, A.V. The need for quantification of uncertainty in artificial intelligence for clinical data analysis: Increasing the level of trust in the decision-making process. *IEEE Syst. Man. Cybern. Mag.* **2022**, *8*, 28–40. [CrossRef]
8. Shrestha, Y.R.; Krishna, V.; von Krogh, G. Augmenting organizational decision-making with deep learning algorithms: Principles, promises, and challenges. *J. Bus. Res.* **2021**, *123*, 588–603. [CrossRef]

9. Kidziński, Ł.; Mohanty, S.P.; Ong, C.F.; Huang, Z.; Zhou, S.; Pechenko, A.; Stelmaszczyk, A.; Jarosik, P.; Pavlov, M.; Kolesnikov, S.; et al. Learning to Run Challenge Solutions: Adapting Reinforcement Learning Methods for Neuromusculoskeletal Environments. In *The NIPS '17 Competition: Building Intelligent Systems*; The Springer Series on Challenges in Machine Learning; Springer: Cham, Switzerland, 2018; pp. 121–153. [\[CrossRef\]](#)
10. Nikolic, B.; Ignjatic, J.; Suzic, N.; Stevanov, B.; Rikalovic, A. Predictive manufacturing systems in industry 4.0: Trends, benefits and challenges. In *Proceedings of the 28th DAAAM International Symposium on Intelligent, Zadar, Croatia, 8–11 November 2017*; pp. 796–0802. [\[CrossRef\]](#)
11. Dogan, A.; Birant, D. Machine learning and data mining in manufacturing. *Expert Syst. Appl.* **2021**, *166*, 114060. [\[CrossRef\]](#)
12. Alzubaidi, L.; Zhang, J.; Humaidi, A.J.; Al-Dujaili, A.; Duan, Y.; Al-Shamma, O.; Santamaría, J.; Fadhel, M.A.; Al-Amidie, M.; Farhan, L. Review of deep learning: Concepts, CNN architectures, challenges, applications, future directions. *J. Big Data* **2021**, *8*, 53. [\[CrossRef\]](#) [\[PubMed\]](#)
13. Jiang, J.; Xiong, Y.; Zhang, Z.; Rosen, D.W. Machine learning integrated design for additive manufacturing. *J. Intell. Manuf.* **2020**, *33*, 1073–1086. [\[CrossRef\]](#)
14. Baumers, M.; Özcan, E. Scope for Machine Learning in Digital Manufacturing. 2016. Available online: <http://tinyurl.com/zfpvzau> (accessed on 12 February 2026).
15. Thompson, M.K.; Moroni, G.; Vaneker, T.; Fadel, G.; Campbell, R.I.; Gibson, I.; Bernard, A.; Schulz, J.; Graf, P.; Ahuja, B.; et al. Design for Additive Manufacturing: Trends, opportunities, considerations, and constraints. *CIRP Ann.* **2016**, *65*, 737–760. [\[CrossRef\]](#)
16. Li, H. Improve the performance of a complex fms with a hybrid machine learning algorithm. *J. Softw. Eng. Appl.* **2017**, *10*, 257–272.
17. Gordon, J.; Sohal, A.S. Assessing manufacturing plant competitiveness—An empirical field study. *Int. J. Oper. Prod. Manag.* **2001**, *21*, 233–253. [\[CrossRef\]](#)
18. Shiang, L.; Nagaraj, S. Impediments to innovation: Evidence from Malaysian manufacturing firms. *Asia Pac. Bus. Rev.* **2011**, *17*, 209–223. [\[CrossRef\]](#)
19. Thomas, A.J.; Byard, P.; Evans, R. Identifying the UK's manufacturing challenges as a benchmark for future growth. *J. Manuf. Technol. Manag.* **2012**, *23*, 142–156. [\[CrossRef\]](#)
20. Wuest, T. Hybrid Analytics in Smart Manufacturing Systems—Addressing the Manufacturing Data Challenge. *Smart Sustain. Manuf. Syst.* **2024**, *8*, 110–118. [\[CrossRef\]](#)
21. Hoang, D.; Wiegatz, K. Machine learning methods in finance: Recent applications and prospects. *Eur. Financ. Manag.* **2023**, *29*, 1657–1701. [\[CrossRef\]](#)
22. Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A.A.; Veness, J.; Bellemare, M.G.; Graves, A.; Riedmiller, M.; Fidjeland, A.K.; Ostrovski, G.; et al. Human-level control through deep reinforcement learning. *Nature* **2015**, *518*, 529–533. [\[CrossRef\]](#) [\[PubMed\]](#)
23. Pouyanfar, S.; Sadiq, S.; Yan, Y.; Tian, H.; Tao, Y.; Reyes, M.P.; Shyu, M.-L.; Chen, S.-C.; Iyengar, S.S. A survey on deep learning: Algorithms, techniques, and applications. *ACM Comput. Surv.* **2018**, *51*, 1–36. [\[CrossRef\]](#)
24. He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; Girshick, R. Masked autoencoders are scalable vision learners. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, New Orleans, LA, USA, 18–24 June 2022; pp. 15979–15988.
25. Zohuri, B. Shielding of Ionizing Radiation in Nuclear Power Reactor Environments (A Technical Memorandum). *J. Mater. Sci. Manuf. Res.* **2022**, *3*, 1–7. [\[CrossRef\]](#)
26. Baduge, S.K.; Thilakarathna, S.; Perera, J.S.; Arashpour, M.; Sharafi, P.; Teodosio, B.; Shringi, A.; Mendis, P. Artificial intelligence and smart vision for building and construction 4.0: Machine and deep learning methods and applications. *Autom. Constr.* **2022**, *141*, 104440. [\[CrossRef\]](#)
27. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of the ICLR 2021-9th International Conference on Learning Representations, Virtual Event, 3–7 May 2021*.
28. Gadepally, V.; Goodwin, J.; Kepner, J.; Reuther, A.; Reynolds, H.; Samsi, S.; Su, J.; Martinez, D. AI Enabling Technologies: A Survey. *arXiv* **2019**, arXiv:1905.03592.
29. Shu, K.; Mahudeswaran, D.; Wang, S.; Lee, D.; Liu, H. FakeNewsNet: A Data Repository with News Content, Social Context, and Spatiotemporal Information for Studying Fake News on Social Media. *Big Data* **2020**, *8*, 171–188. [\[CrossRef\]](#) [\[PubMed\]](#)
30. Dastres, R.; Soori, M. Artificial neural network systems. *Int. J. Imaging Robot.* **2021**, *21*, 13–25.
31. Karoly, A.I.; Galambos, P.; Kuti, J.; Rudas, I.J. Deep learning in robotics: Survey on model structures and training strategies. *IEEE Trans. Syst. Man. Cybern. Syst.* **2020**, *51*, 266–279. [\[CrossRef\]](#)
32. Zhao, W.; Queralta, J.P.; Westerlund, T. Sim-to-real transfer in deep reinforcement learning for robotics: A survey. In *Proceedings of the 2020 IEEE Symposium Series on Computational Intelligence (SSCI)*, Canberra, ACT, Australia, 1–4 December 2020; pp. 737–744. [\[CrossRef\]](#)

33. Kapeller, A.; Felzmann, H.; Fosch-Villaronga, E.; Hughes, A.-M. A Taxonomy of Ethical, Legal and Social Implications of Wearable Robots: An Expert Perspective. *Sci. Eng. Ethics* **2020**, *26*, 3229–3247. [\[CrossRef\]](#) [\[PubMed\]](#)
34. Dimiduk, D.M.; Holm, E.A.; Niezgoda, S.R. Perspectives on the Impact of Machine Learning, Deep Learning, and Artificial Intelligence on Materials, Processes, and Structures Engineering. *Integr. Mater. Manuf. Innov.* **2018**, *7*, 157–172. [\[CrossRef\]](#)
35. Tyagi, A.K.; Chahal, P. Artificial Intelligence and Machine Learning Algorithms. In *Research Anthology on Machine Learning Techniques, Methods, and Applications*; Information Resources Management Association, Ed.; IGI Global: Hershey, PA, USA, 2022; pp. 421–446. [\[CrossRef\]](#)
36. Linaza, M.T.; Posada, J.; Bund, J.; Eisert, P.; Quartulli, M.; Döllner, J.; Pagani, A.; Olaizola, I.G.; Barriguinha, A.; Moysiadis, T.; et al. Data-Driven Artificial Intelligence Applications for Sustainable Precision Agriculture. *Agronomy* **2021**, *11*, 1227. [\[CrossRef\]](#)
37. Garg, P.K. Overview of Artificial Intelligence. In *Artificial Intelligence: Technologies, Applications, and Challenges*; Sharma, L., Garg, P.K., Eds.; Chapman and Hall/CRC: New York, NY, USA, 2021; pp. 3–18. [\[CrossRef\]](#)
38. Kamruzzaman, M.M. New opportunities, challenges, and applications of edge-AI for connected healthcare in smart cities. In Proceedings of the 2021 IEEE Globecom Workshops (GC Wkshps), Madrid, Spain, 7–11 December 2021. Available online: <https://ieeexplore.ieee.org/abstract/document/9682055/> (accessed on 28 April 2025).
39. Alanazi, A. Using machine learning for healthcare challenges and opportunities. *Inform. Med. Unlocked* **2022**, *30*, 100924. [\[CrossRef\]](#)
40. Biliiani, I.; Skamnia, E.; Economou, P.; Zacharias, I. A Novel Methodology to Correct Chlorophyll-a Concentrations from Satellite Data and Assess Credible Phenological Patterns. *Remote Sens.* **2025**, *17*, 1156. [\[CrossRef\]](#)
41. Silver, D.; Schrittwieser, J.; Simonyan, K.; Antonoglou, I.; Huang, A.; Guez, A.; Hubert, T.; Baker, L.; Lai, M.; Bolton, A.; et al. Mastering the game of Go without human knowledge. *Nature* **2017**, *550*, 354–359. [\[CrossRef\]](#) [\[PubMed\]](#)
42. Dalzochio, J.; Kunst, R.; Pignaton, E.; Binotto, A.; Sanyal, S.; Favilla, J.; Barbosa, J. Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges. *Comput. Ind.* **2020**, *123*, 103298. [\[CrossRef\]](#)
43. Azari, M.S.; Flammini, F.; Santini, S.; Caporuscio, M. A systematic literature review on transfer learning for predictive maintenance in industry 4.0. *IEEE Access* **2023**, *11*, 12887–12910. [\[CrossRef\]](#)
44. Ali, S.; Abuhmed, T.; El-Sappagh, S.; Muhammad, K.; Alonso-Moral, J.M.; Confalonieri, R.; Guidotti, R.; Del Ser, J.; Díaz-Rodríguez, N.; Herrera, F. Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence. *Inf. Fusion* **2023**, *99*, 101805. [\[CrossRef\]](#)
45. Lundberg, S.M.; Lee, S.-I. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems 30*; Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R., Eds.; Curran Associates, Inc.: Red Hook, NY, USA, 2017. Available online: <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html> (accessed on 30 May 2026).
46. Minh, D.L.; Sadeghi-Niaraki, A.; Huy, H.D.; Min, K.; Moon, H. Deep learning approach for short-term stock trends prediction based on two-stream gated recurrent unit net-work. *IEEE Access* **2018**, *6*, 55392–55404. [\[CrossRef\]](#)
47. Santos, M.S.; Soares, J.P.; Abreu, P.H.; Araujo, H.; Santos, J. Cross-validation for imbalanced datasets: Avoiding overoptimistic and overfitting approaches. *IEEE Comput. Intell. Mag.* **2018**, *13*, 59–76. [\[CrossRef\]](#)
48. Mahdavinnejad, M.S.; Rezvan, M.; Barekatin, M.; Adibi, P.; Barnaghi, P.; Sheth, A.P. Machine learning for internet of things data analysis: A survey. *Digit. Commun. Netw.* **2018**, *4*, 161–175. [\[CrossRef\]](#)
49. Chang, Z.; Lei, L.; Zhou, Z.; Mao, S.; Ristaniemi, T. Learn to cache: Machine learning for network edge caching in the big data era. *IEEE Wirel. Commun.* **2018**, *25*, 28–35. [\[CrossRef\]](#)
50. Teo, Z.L.; Jin, L.; Liu, N.; Li, S.; Miao, D.; Zhang, X.; Ng, W.Y.; Tan, T.F.; Lee, D.M.; Chua, K.J.; et al. Federated machine learning in healthcare: A systematic review on clinical applications and technical architecture. *Cell Rep. Med.* **2024**, *5*, 101419. [\[CrossRef\]](#) [\[PubMed\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.