Proceedings of CONF-MLA 2026 Symposium: Intelligent Systems and Automation: AI Models, IoT, and Robotic Algorithms
DOI: 10.54254/2755-2721/2026.LD38249

The Applications and Challenges of Machine Learning in Intelligent Decision-Making

Chengkai Wu

Faculty of Innovation Engineering, Macau University of Science and Technology, Macau, China
wuchengkai070214@gmail.com

Abstract. Machine learning has become a central technology for transforming data into intelligent decisions in modern organisations. Recent advances in deep learning and reinforcement learning have greatly expanded the capacity of machines to recognise patterns, forecast outcomes and optimise sequential choices, which has stimulated a wide range of decision-support applications. By adopting a structured review method, this paper reviews the applications and challenges of machine learning in intelligent decision-making. The results show that learning-based models improve the accuracy, speed and scalability of decisions in clinical diagnosis, asset management, supply-chain planning and autonomous driving. The analysis also identified four recurring challenges that limit practical adoption, namely data quality, model interpretability, algorithmic fairness and safety. The paper concluded that reliable deployment depends on transparent model design, fairness auditing and human-in-the-loop governance, and it proposed directions for future research in hybrid human-machine decision systems.

Keywords: machine learning, intelligent decision-making, deep learning, reinforcement learning, interpretability

1. Introduction

Machine learning has transformed the way organisations convert data into decisions. Advances in deep neural networks have produced remarkable results in image recognition, natural language processing and strategic game playing [1, 2]. In parallel, reinforcement learning has demonstrated the ability to master sequential decision tasks, from playing Atari video games to defeating human champions in the game of Go [3, 4]. These successes have stimulated growing interest in intelligent decision-making, in which algorithms assist or replace humans in choosing among alternative actions under uncertainty. Both public institutions and private firms now rely on such systems for tasks ranging from clinical screening to investment and logistics planning.

Despite this rapid progress, the research field remains fragmented. Most studies concentrate on a single domain, such as medicine or finance, and few provide a cross-domain account of how machine learning supports decisions and where its limits lie. Important questions concern the conditions under which learning-based models outperform traditional decision tools, the obstacles that prevent routine deployment, and the ways in which machines and humans should share responsibility for consequential choices.

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2026.LD38249

This paper reviews the applications and challenges of machine learning in intelligent decision-making across four decision domains, namely healthcare, finance, operations management and transportation. It employs a structured method to analyze peer-reviewed studies based on decision task, learning paradigm and reported outcome. It contributes to the design of decision systems that are not only accurate but also trustworthy and accountable.

2. Machine learning paradigms and their roles in decision-making

2.1. Core learning paradigms

Machine learning methods applied to decision problems can be grouped into three paradigms. Supervised learning constructs a mapping from labelled input–output pairs and is used when historical data contain explicit outcomes. In decision contexts, supervised models generate predictions or classifications, such as the probability of disease, the creditworthiness of a borrower or the expected demand for a product, which then enter the decision process. Unsupervised learning operates on unlabelled data and seeks hidden structure, including clusters, associations and anomalies. It facilitates decision-making by segmenting customers, detecting fraudulent transactions and identifying unusual operating states. Reinforcement learning differs from the other two paradigms because an agent learns a policy through trial and error while receiving rewards from an environment, and it is therefore naturally suited to sequential decisions in which current actions affect future states [5]. These three paradigms are rarely used in isolation, and many practical decision systems combine them with statistical inference and domain knowledge. Table 1 summarises the paradigms and their typical decision tasks.

Table 1. Comparison of machine learning paradigms in decision-making

Paradigm	Learning signal	Decision-related objective	Typical decision task
Supervised learning	Labelled historical examples	Accurately predict an outcome	Medical diagnosis, credit scoring, demand forecasting
Unsupervised learning	Unlabelled examples	Discover latent structure	Customer segmentation, fraud detection, anomaly monitoring
Reinforcement learning	Rewards from the environment	Maximise cumulative reward over time	Game playing, dynamic pricing, robotic control

2.2. From prediction to decision: a general framework

In practice, machine learning supports intelligent decisions through a pipeline rather than through a single model. The process begins with data collection and cleaning, followed by feature construction and model training with validation. The trained model then produces predictions, probabilities or recommended actions, and a separate decision layer combines these outputs with business rules, costs and risk constraints before a final action is executed. The decision layer is essential because a prediction is rarely the decision itself, for example, a forecast of demand must be translated into ordering quantities under storage and budget limits, and a diagnosis must be converted into a treatment plan by a clinician. A feedback loop closes the pipeline, because new outcomes are recorded and used to retrain the models. This framework implies that decision quality depends not only on model accuracy but also on data representativeness, uncertainty quantification and the design of the surrounding decision rules. Consequently, the successful application of machine

DOI: 10.54254/2755-2721/2026.LD38249

learning requires cooperation between data scientists, domain experts and decision-makers throughout the entire pipeline.

3. Applications of machine learning in intelligent decision-making

Machine learning has been applied to decisions in almost every sector of the economy.

3.1. Clinical decision support in healthcare

In healthcare, machine learning assists decisions ranging from screening and diagnosis to treatment selection and hospital management. Esteva et al. trained a deep convolutional neural network on more than 129,000 clinical images covering over 2,000 skin diseases and demonstrated that the model classified skin lesions with an accuracy comparable to that of board-certified dermatologists [6]. Such systems can support clinicians in early screening, triage and workload reduction. Topol characterised this development as the convergence of human and artificial intelligence in high-performance medicine, in which machines perform routine pattern recognition while clinicians supply contextual understanding, empathy and final judgement [7]. Davenport and Kalakota reviewed the broader potential of artificial intelligence in healthcare, including diagnostic support, patient monitoring and administrative optimisation, and they emphasised that issues of liability, data privacy and regulation remain unresolved before such systems can be deployed routinely [8]. The reviewed evidence indicates that machine learning is best framed as a decision aid rather than as a replacement for clinical reasoning.

3.2. Investment and risk decisions in finance

In finance, machine learning supports investment, credit and risk decisions under conditions of high dimensionality and uncertainty. Gu, Kelly and Xiu compared a broad set of machine learning methods, including linear models, penalised regressions, tree ensembles and neural networks, for predicting the cross-section of United States stock returns [9]. The study showed that tree-based methods and neural networks produce the most accurate return forecasts and yield portfolios with substantially higher out-of-sample performance than traditional linear factor models. These results illustrate how machine learning can capture nonlinear and interactive relationships among firm characteristics that conventional models overlook. Similar techniques are applied to credit scoring, fraud detection and algorithmic trading, where the ability to process large and heterogeneous datasets provides a clear advantage, although interpretability and overfitting remain important concerns in regulated financial decisions.

3.3. Operations and supply chain management

In operations and supply chain management, machine learning improves demand forecasting, inventory control, capacity planning and risk management. Choi, Wallace and Wang reviewed the use of big data analytics in operations management and demonstrated that learning-based forecasting methods consistently reduce prediction error and allow more responsive replenishment decisions than conventional time-series approaches [10]. Their review also showed that the full benefits of machine learning are realised only when predictive models are integrated with optimisation models and when data quality, privacy and organisational capability are managed jointly. In manufacturing and retail, these capabilities enable dynamic pricing, predictive maintenance and the early detection

DOI: 10.54254/2755-2721/2026.LD38249

of supply-chain disruptions, all of which require decisions to be updated frequently as new data arrive.

3.4. Autonomous driving and transportation

Autonomous driving represents one of the most demanding decision environments for machine learning, because vehicles must make rapid and safe decisions in continuous, safety-critical settings. Yurtsever, Lambert, Carballo and Takeda surveyed the common practices and emerging technologies of autonomous driving, covering environment perception, localisation, motion planning and vehicle control [11]. The survey showed that deep learning dominates perception tasks such as object detection and semantic segmentation, whereas reinforcement learning and optimisation are used for planning and control. Decisions in autonomous vehicles occur at multiple levels, from strategic route selection to millisecond-level braking choices, and the greatest obstacle remains the validation of safety across rare and unpredictable traffic situations.

4. Challenges in the adoption of machine learning for intelligent decision-making

4.1. Data quality and availability

All learning-based decision systems depend on the quality of the data from which they are trained. Noise, missing values, class imbalance and errors in historical labels can propagate silently into decisions, and models trained on past data often degrade when the underlying distribution shifts over time. In high-stakes domains, relevant data are frequently fragmented across institutions and restricted by privacy regulations, which makes it difficult to assemble representative training sets. In healthcare, for example, diagnostic datasets vary in acquisition devices and patient populations, so models may underperform when applied to new hospitals [7]. In finance, structural changes in markets can invalidate relationships learned from historical data [9]. These issues imply that data governance, continuous monitoring and mechanisms for updating models are prerequisites for reliable deployment.

4.2. Interpretability and trust

A second challenge concerns interpretability. Deep learning models are often opaque, and decision-makers, regulators and affected individuals increasingly demand explanations for consequential outcomes such as loan rejections or medical recommendations. Rudin argued that for high-stakes decisions, inherently interpretable models should be preferred over black-box models with post-hoc explanations, because post-hoc explanations can be inaccurate and misleading [12]. Ribeiro, et al. introduced LIME, a method that explains the predictions of any classifier by fitting an interpretable model locally around a given instance, thereby allowing users to judge whether a prediction can be trusted [13]. The literature distinguishes between global interpretability, which describes how a model behaves in general, and local interpretability, which explains individual predictions. Since accuracy and interpretability frequently involve a trade-off, the design of decision systems must match model transparency to the stakes and accountability requirements of the task.

4.3. Fairness and bias

Machine learning systems can also reproduce or amplify social biases, which raises serious concerns in decisions that affect employment, credit, criminal justice and healthcare. Mehrabi, Morstatter,

DOI: 10.54254/2755-2721/2026.LD38249

Saxena, Lerman and Galstyan provided a comprehensive survey of bias and fairness in machine learning, categorising the sources of bias that arise during data generation, model training and deployment, and reviewing mitigation approaches [14]. Biases can enter through unrepresentative samples, prejudiced labels or feedback loops in which model outputs influence future data. Their consequences include disparate treatment of protected groups and the reinforcement of existing inequalities. Mitigation requires a combination of fair data collection, fairness metrics, systematic auditing before deployment and transparent documentation, together with regulatory oversight, because fairness is a sociotechnical property that cannot be guaranteed by algorithms alone.

4.4. Safety, robustness and human oversight

Safety and robustness constitute the fourth challenge. Models can be misled by small adversarial perturbations, and reinforcement learning agents may exploit unintended reward signals, producing behaviour that appears successful during training but is unsafe in practice. In autonomous driving, rare edge cases are difficult to anticipate during development, and the consequences of failure are severe [11]. At the same time, human interaction with machine recommendations raises the problem of automation bias, in which decision-makers over-trust algorithmic outputs, or the opposite risk of under-using valuable assistance. Jarrahi argued that effective decision-making requires a symbiosis between humans and artificial intelligence, in which algorithms contribute scale, speed and pattern recognition while humans contribute context, ethical judgement and accountability, particularly in ambiguous and novel situations [15]. Governance structures, human-in-the-loop protocols and the clear allocation of responsibility are therefore necessary components of intelligent decision systems.

5. Conclusion

This paper reviews the applications and challenges of machine learning in intelligent decision-making. The review showed that machine learning has become a powerful component of decision systems in healthcare, finance, operations management and transportation. Supervised learning provides accurate predictions and classifications that inform routine decisions, unsupervised learning reveals structure that supports segmentation and anomaly detection, and reinforcement learning enables sequential optimisation in dynamic environments. Representative applications demonstrate measurable benefits: deep convolutional networks classify skin lesions at the level of expert dermatologists, machine learning models improve out-of-sample asset returns, learning-based forecasts reduce supply-chain costs, and autonomous driving systems integrate perception and planning through deep learning.

At the same time, the analysis identified four recurring challenges that constrain widespread adoption. First, decision quality depends critically on data representativeness, quality and continuity, and fragmented or biased data can invalidate model outputs. Second, the opacity of many models conflicts with the accountability required in high-stakes decisions, so interpretable design and local explanation methods are needed. Third, learning-based systems can amplify historical bias and cause unfair outcomes, which requires fairness auditing and governance. Fourth, safety and robustness concerns, together with the risks of automation bias, demand human oversight and the clear allocation of responsibility.

The paper therefore concludes that intelligent decision-making should not be viewed as the replacement of human judgement by algorithms, but as the design of hybrid systems in which machines and humans complement each other. Several directions for future research follow from this conclusion. Future work should develop interpretable and causally informed models that remain

DOI: 10.54254/2755-2721/2026.LD38249

accurate at scale, establish standardised frameworks for auditing fairness and robustness, and quantify uncertainty so that decision-makers can act on model confidence rather than on point predictions alone. Organisations should invest in data governance, continuous model monitoring and interdisciplinary teams that combine technical and domain expertise. Because the reviewed evidence is heterogeneous and the field evolves rapidly, the conclusions of this review should be interpreted as a snapshot of current practice. Future reviews may benefit from systematic meta-analysis as more deployment studies become available. Overall, machine learning offers substantial opportunities for more accurate, faster and better-informed decisions, and the realisation of these opportunities will depend on addressing the technical, organisational and ethical challenges identified in this paper.

References

[1] LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature, 521*, 436–444.
[2] Jordan, M. I., & Mitchell, T. M. (2015). Machine learning: Trends, perspectives, and prospects. *Science, 349*, 255–260.
[3] Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature, 529*, 484–489.
[4] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature, 518*, 529–533.
[5] Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.
[6] Esteva, A., Kuprel, B., Novoa, R. A., Ko, J., Swetter, S. M., Blau, H. M., & Thrun, S. (2017). Dermatologist-level classification of skin cancer with deep neural networks. *Nature, 542*, 115–118.
[7] Topol, E. J. (2019). High-performance medicine: The convergence of human and artificial intelligence. *Nature Medicine, 25*, 44–56.
[8] Davenport, T., & Kalakota, R. (2019). The potential for artificial intelligence in healthcare. *Future Healthcare Journal, 6*, 94–98.
[9] Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies, 33*(5), 2223–2273.
[10] Choi, T.-M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. *Production and Operations Management, 27*(10), 1868–1883.
[11] Yurtsever, E., Lambert, J., Carballo, A., & Takeda, K. (2020). A survey of autonomous driving: Common practices and emerging technologies. *IEEE Access, 8*, 58443–58469.
[12] Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*, 206–215.
[13] Ribeiro, M. T., Singh, S., & Guestrin, C. (2016, August 13-17). "Why should I trust you?": Explaining the predictions of any classifier. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135-1144. ACM.
[14] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys, 54*(6), 1-35.
[15] Jarrahi, M. H. (2018). Artificial intelligence and the future of work: Human-AI symbiosis in organizational decision making. *Business Horizons, 61*(4), 577-586.