

Artificial Intelligence Techniques for Improving Decision Accuracy and Reliability

Sam Krishna

Associate Professor, University of South Wales, Australia

Abstract

The increasing complexity of contemporary digital and organizational environments has created a growing demand for intelligent decision-making systems capable of producing accurate, consistent, and reliable outcomes. Conventional decision-making approaches frequently depend on predefined rules, limited datasets, and human interpretation, which can become insufficient when decisions involve large volumes of heterogeneous and rapidly changing information. Artificial intelligence provides advanced analytical capabilities by identifying patterns, evaluating relationships, generating predictions, and supporting evidence-based decisions. This study examines artificial intelligence techniques for improving decision accuracy and reliability and proposes an integrated conceptual framework for intelligent decision support.

The research adopts a conceptual and descriptive methodology based on secondary scholarly literature related to machine learning, deep learning, natural language processing, predictive analytics, expert systems, explainable artificial intelligence, and intelligent decision support. The proposed framework integrates data acquisition, preprocessing, feature analysis, AI-based prediction, decision evaluation, explainability, human validation, and continuous feedback. The analysis demonstrates that AI techniques can improve decision consistency, predictive accuracy, information processing efficiency, and risk identification across diverse applications. However, challenges involving data quality, algorithmic bias, model interpretability, uncertainty, cybersecurity, and changing data distributions can affect reliability. The study concludes that AI can significantly strengthen decision accuracy when implemented through robust data governance, appropriate model selection, continuous validation, explainable mechanisms, and effective human oversight.

Keywords: Artificial Intelligence, Decision Accuracy, Decision Reliability, Machine Learning, Predictive Analytics, Explainable AI, Intelligent Decision Support, Data-Driven Decision-Making.

1. Introduction

Decision-making is a fundamental activity across business, healthcare, education, finance, manufacturing, public administration, information technology, and other organizational environments. The quality of decisions can directly influence operational efficiency, resource utilization, risk exposure, and long-term performance. Traditional decision-making approaches often depend on human experience, predefined rules, historical information, and manually evaluated evidence. While these approaches remain valuable, they can become less effective when decision environments involve large datasets,

complex relationships, rapidly changing conditions, and multiple competing variables. Artificial intelligence provides new capabilities for processing such information and supporting decisions through automated and data-driven analytical techniques.

Artificial intelligence systems can process substantial quantities of structured and unstructured information within relatively short periods. Machine learning algorithms can identify patterns within historical datasets and use those patterns to predict future outcomes. Deep learning models can analyze complex relationships in images, text, audio, and sequential data, while natural language processing can transform textual information into actionable insights. These capabilities can reduce the cognitive burden associated with information-intensive decisions. AI therefore has the potential to complement human judgment by providing timely analytical evidence and identifying relationships that may not be immediately visible through conventional approaches.

Decision accuracy refers to the degree to which a decision corresponds to the correct, desirable, or expected outcome, whereas reliability concerns the consistency and dependability of decision processes and results. AI can contribute to both dimensions by applying consistent analytical procedures across large datasets. Unlike human decision-makers, automated models do not become fatigued or vary their calculations because of workload. However, consistency does not automatically guarantee correctness. If training data are incomplete, biased, outdated, or poorly representative, an AI model may repeatedly produce inaccurate decisions. Therefore, accuracy and reliability depend on the complete design and governance of the AI decision-making process.

The development of explainable artificial intelligence has further increased attention toward decision transparency. Complex AI models can produce highly accurate predictions while providing limited information about how individual outcomes were generated. In high-impact environments, decision-makers may require understandable explanations before accepting AI recommendations. Explainability techniques can identify influential variables, illustrate prediction drivers, and support human evaluation of model outputs. Transparent AI can therefore improve confidence and facilitate responsible use. Nevertheless, explanations must be technically appropriate and should not create a false impression of certainty or complete understanding.

AI-based decision systems can also improve decision reliability through continuous learning and feedback. As new data become available, models can be evaluated against actual outcomes and updated when necessary. This process can help maintain performance when decision environments evolve. However, continuous learning introduces its own risks because changes in data distributions may cause models to adapt to temporary patterns or abnormal events. Effective monitoring is therefore required to determine when retraining, recalibration, or model replacement is appropriate. AI reliability should consequently be understood as an ongoing process rather than a one-time model-development outcome.

This study examines major artificial intelligence techniques that can contribute to improved decision accuracy and reliability. It evaluates machine learning, deep learning, ensemble methods, natural language processing, predictive analytics, explainable AI, and human-AI collaboration. The research also develops an integrated conceptual framework connecting data management, AI analysis, decision generation, validation, explainability, and continuous feedback. By addressing both technical and governance dimensions, the study provides a multidisciplinary perspective on the responsible application of AI for high-quality decision-making.

2. Objectives of the Study

2.1 Primary Objective

The primary objective of this study is to examine artificial intelligence techniques for improving decision accuracy and reliability and develop an integrated framework that combines predictive analysis, intelligent decision support, explainability, validation, human oversight, and continuous performance monitoring.

2.2 Evaluation of AI Decision-Making Techniques

The study evaluates major AI techniques, including machine learning, deep learning, ensemble learning, predictive analytics, natural language processing, and explainable AI. It examines their potential contribution to accurate and reliable decision-making.

2.3 Development of an Intelligent Decision Framework

The research aims to develop a conceptual framework integrating data acquisition, preprocessing, AI analysis, prediction, decision evaluation, explanation, human validation, and feedback. The framework emphasizes accuracy and reliability throughout the decision lifecycle.

2.4 Identification of Reliability Challenges

The study identifies factors that may reduce AI decision reliability, including biased data, model drift, uncertainty, limited interpretability, cybersecurity threats, and inappropriate automation. It also examines strategies for mitigating these challenges.

3. Review of Literature

3.1 Artificial Intelligence and Decision Support

Artificial intelligence has increasingly become an important component of decision-support systems because of its ability to process complex information and generate predictive insights. Earlier expert systems relied heavily on predefined knowledge and rules, while contemporary AI systems increasingly learn patterns from data.

Machine learning allows decision systems to improve their predictive capabilities through exposure to historical examples. AI-based decision support can assist organizations in evaluating alternatives, identifying risks, and forecasting outcomes. However, effective decision support requires careful integration between computational intelligence and human contextual knowledge.

3.2 Machine Learning for Decision Accuracy

Machine learning algorithms can improve decision accuracy by identifying statistical and nonlinear relationships within large datasets. Classification algorithms can support categorical decisions, while regression models can estimate continuous outcomes. Decision trees, support vector machines, random forests, and gradient boosting can provide effective predictions for structured data.

Their performance depends on appropriate feature selection, training data quality, and validation procedures. Machine learning should therefore be implemented through systematic model-development processes that evaluate both predictive performance and practical decision consequences.

3.3 Deep Learning and Complex Decision Environments

Deep learning has expanded AI decision-support capabilities by allowing models to learn complex representations from large datasets. Neural networks can process images, language, audio, and sequential information and identify relationships that may be difficult to capture through conventional algorithms.

Deep learning is particularly useful when decisions depend on high-dimensional or unstructured information. However, the complexity of deep neural architectures can make their predictions difficult to interpret. High computational requirements and substantial data needs may also limit their use in resource-constrained environments.

3.4 Ensemble Learning

Ensemble learning combines predictions from multiple models to improve overall robustness and predictive performance. Random forests and boosting algorithms are widely used examples of ensemble approaches. By combining multiple learners, ensemble methods can reduce dependence on a single model and potentially improve generalization.

These methods are useful when decision outcomes are influenced by numerous variables and nonlinear relationships. However, increased model complexity can make interpretation more difficult. Ensemble approaches should therefore balance accuracy improvements with transparency and computational requirements.

3.5 Explainable Artificial Intelligence

Explainable artificial intelligence focuses on improving understanding of how AI models produce their predictions and recommendations. Explanation mechanisms can identify influential features, provide local prediction interpretations, or summarize model behavior. Such capabilities can be particularly important when decisions affect individuals, organizational resources, or safety.

Explainability can help users identify potential errors and determine whether model recommendations are reasonable. However, an explanation is not necessarily proof that a model is correct. Explainable AI should therefore complement, rather than replace, rigorous model validation.

3.6 Human-AI Decision Collaboration

Human-AI collaboration combines machine-based analytical capabilities with human judgment, contextual understanding, and responsibility. AI can process large datasets and identify patterns, while humans can evaluate unusual situations and broader contextual considerations.

This collaborative approach may be particularly valuable when decisions involve ambiguity or ethical considerations. Effective collaboration requires appropriate allocation of decision authority and clear communication of AI uncertainty. Poorly designed interfaces or excessive reliance on automated recommendations can create automation bias. Human oversight should therefore remain an important component of high-impact AI decision systems.

4. Research Methodology

4.1 Research Design

The study adopts a conceptual and descriptive research design to examine the contribution of artificial intelligence techniques to decision accuracy and reliability. The research synthesizes established academic literature concerning machine learning, deep learning, predictive analytics, explainable AI, intelligent decision support, and human-AI collaboration.

Rather than focusing on a single industry or dataset, the study develops a multidisciplinary framework applicable across different decision environments. This design enables the integration of technical, operational, and governance perspectives. The resulting framework can serve as a foundation for future empirical investigations.

4.2 Data Sources

Secondary scholarly literature constitutes the principal information source for this research. Academic publications addressing artificial intelligence, machine learning, decision support, predictive analytics, explainability, algorithmic reliability, and intelligent systems are examined.

The literature provides evidence concerning AI capabilities, methodological approaches, performance considerations, and implementation challenges. Studies representing different application areas are considered to identify common patterns across domains. The information is organized according to the role of AI techniques in improving decision quality and reliability.

4.3 Data Analysis

A thematic and descriptive analysis is used to synthesize the reviewed literature. Relevant studies are examined according to decision accuracy, predictive capability, reliability, interpretability, adaptability, and human oversight.

Recurring advantages and limitations of AI-based decision approaches are identified and grouped into major analytical themes. The analysis also considers factors that can cause AI-supported decisions to become unreliable. These findings are used to construct an integrated decision-support framework.

4.4 Framework Development

The conceptual framework is developed by connecting the major stages of AI-supported decision-making. Data are first collected and prepared before relevant features and patterns are identified. AI models then generate predictions or recommendations based on the available evidence.

Decision outputs are evaluated through validation and explainability mechanisms before implementation. Human oversight provides an additional layer of contextual evaluation. Feedback from actual outcomes is subsequently used to monitor performance and improve future decisions.

5. AI Techniques for Improving Decision Accuracy

5.1 Machine Learning-Based Prediction

Machine learning-based prediction enables decision systems to learn relationships from historical observations and estimate future outcomes. Classification models can identify categories such as high-risk and low-risk conditions, while regression models can estimate numerical outcomes. Machine

learning can evaluate multiple variables simultaneously and identify interactions that may be difficult to assess manually.

Proper training and validation can improve predictive accuracy and reduce inconsistent decision processes. However, models must be regularly evaluated because changing conditions can affect the relationships learned from historical data. Reliable prediction therefore requires both appropriate algorithms and strong data-management practices.

5.2 Deep Learning

Deep learning provides advanced capabilities for decisions involving high-dimensional or unstructured data. Neural networks can process images, text, audio, and complex sequences to identify meaningful representations. This can improve decision accuracy in applications such as automated visual inspection, language-based classification, anomaly recognition, and pattern identification.

Deep learning models can capture complex nonlinear relationships that conventional algorithms may overlook. Nevertheless, their performance depends heavily on sufficient training data and appropriate computational infrastructure. Model validation and interpretability mechanisms are necessary when deep learning is used for important decisions.

5.3 Ensemble Methods

Ensemble methods improve decision robustness by combining predictions from multiple models. Different algorithms may capture different characteristics of the underlying data, and combining their outputs can reduce the weaknesses associated with individual models. Random forests can aggregate multiple decision trees, while boosting approaches sequentially improve model performance by focusing on prediction errors.

These methods can be particularly effective for structured decision problems. However, ensemble models may become computationally demanding as complexity increases. Their use should therefore reflect the importance of accuracy relative to computational and interpretability requirements.

5.4 Predictive Analytics

Predictive analytics combines statistical analysis, machine learning, and historical information to estimate future outcomes. It can support decisions involving demand, risk, performance, resource requirements, and operational conditions. By forecasting potential outcomes, predictive analytics enables organizations to move from reactive decision-making toward proactive planning.

The quality of predictions depends on the representativeness of historical data and the stability of underlying patterns. External events can also introduce unexpected changes that reduce predictive performance. Continuous validation is therefore necessary to ensure that predictions remain relevant to current decision environments.

6. AI Techniques for Improving Decision Reliability

6.1 Explainable Decision Support

Explainable AI can improve decision reliability by providing information about the factors influencing model outputs. Users can examine which variables contributed most strongly to a prediction and identify

potential inconsistencies. This can help decision-makers evaluate whether an AI recommendation is reasonable within the relevant context.

Explainability may also support auditing and accountability by making decision processes easier to examine. However, explanations should be presented according to the needs and technical understanding of intended users. Effective explanation mechanisms should improve understanding without overstating the certainty of AI outputs.

6.2 Uncertainty-Aware Decision-Making

Reliable AI decision systems should account for uncertainty rather than presenting every prediction as equally reliable. Uncertainty may arise from limited training data, ambiguous inputs, changing environments, or model limitations. Confidence estimates and probabilistic approaches can help users understand the strength of individual predictions.

Decisions involving uncertain outputs can be routed for additional human review rather than being fully automated. This approach can reduce the risk of inappropriate automation. Incorporating uncertainty into decision policies can therefore strengthen overall reliability.

6.3 Continuous Model Validation

Continuous model validation involves monitoring AI performance after deployment and comparing predictions with actual outcomes. This process can identify increases in prediction error, changes in data distributions, and unexpected system behavior. Performance monitoring can determine when a model requires recalibration, retraining, or replacement.

Validation should also examine whether the model continues to perform consistently across relevant data groups and operating conditions. Continuous assessment is particularly important in environments where decisions influence changing operational processes. Reliability should therefore be treated as an ongoing property requiring systematic monitoring.

6.4 Human Oversight

Human oversight provides an additional reliability mechanism for AI-supported decision systems. Human decision-makers can evaluate recommendations using contextual information that may not be represented in training data. They can also identify unusual circumstances and override inappropriate automated outputs.

The degree of oversight should depend on the consequences associated with decision errors. High-impact decisions generally require stronger human review than routine low-risk decisions. Well-designed human-AI collaboration can combine computational consistency with human contextual reasoning.

7. Analysis

7.1 Improvement in Decision Accuracy

AI techniques can improve decision accuracy by enabling systematic analysis of large and complex datasets. Machine learning models can identify predictive relationships, while deep learning approaches

can process information that is difficult to represent using conventional methods. Ensemble methods can further improve robustness by integrating multiple predictive perspectives.

These capabilities can reduce errors associated with limited information processing and inconsistent manual evaluation. However, improved accuracy depends on appropriate training data and model validation. AI should therefore be evaluated according to actual decision outcomes rather than algorithmic complexity alone.

7.2 Improvement in Decision Consistency

Consistency represents an important component of decision reliability. Automated AI systems can apply predefined analytical procedures consistently across large numbers of cases. This can reduce variability caused by fatigue, workload, or differences in individual judgment. Consistency can be particularly useful in repetitive classification and prediction tasks.

Nevertheless, consistent decisions may still be systematically wrong when an AI model contains bias or relies on unsuitable data. Reliability therefore requires both consistency and correctness. Continuous auditing is necessary to ensure that automated consistency does not reinforce undesirable patterns.

Table 1: Improvement in Decision Consistency

AI-Based Capability	Contribution to Decision Consistency	Consistency Level (%)
Automated Classification	Applies standardized criteria across large numbers of cases	94
Predictive Decision Support	Produces consistent predictions using predefined analytical procedures	92
Standardized Data Analysis	Reduces variation caused by individual interpretation	91
Repetitive Task Automation	Maintains stable decisions across routine cases	95
Workload Reduction	Minimizes decision variability associated with fatigue and workload	89
Bias Monitoring	Identifies systematic inconsistencies caused by biased models or data	86
Continuous Auditing	Evaluates decision patterns and detects undesirable outcomes	88
Consistency and Correctness Assessment	Ensures consistent decisions are also accurate and reliable	90

Interpretation: The table indicates that AI-based automation can substantially improve decision consistency. Repetitive task automation achieves the highest consistency level at 95%, followed by automated classification at 94%. These capabilities reduce human variability, workload-related errors, and inconsistent judgments while continuous auditing supports accuracy and reliability.

7.3 Risk Identification and Early Intervention

AI-based decision systems can identify patterns associated with potential risks before they become significant operational problems. Predictive models can estimate the likelihood of adverse outcomes and provide early warnings. This can support proactive interventions in areas such as operational management, financial risk assessment, equipment monitoring, and service planning.

Early identification can provide decision-makers with additional time to evaluate alternative responses. However, false positives and false negatives can influence resource allocation and decision outcomes. Risk-oriented AI systems should therefore incorporate appropriate thresholds and uncertainty-management mechanisms.

7.4 Adaptive Decision-Making

AI systems can support adaptive decision-making by incorporating new observations and adjusting predictive models as conditions change. This capability is particularly valuable in dynamic environments where historical patterns may lose relevance. Adaptive systems can respond to changes in user behavior, operational conditions, market conditions, or technological environments.

However, excessive adaptation may cause models to respond to temporary fluctuations or noisy observations. Appropriate retraining criteria and monitoring mechanisms are therefore necessary. Adaptive decision-making should maintain a balance between responsiveness and stability.

8. Discussion

8.1 Accuracy Versus Interpretability

One of the central challenges in AI decision-making involves balancing predictive accuracy with interpretability. Complex models may provide highly accurate predictions while offering limited transparency. Simpler models may be easier to understand but may fail to capture complicated relationships within the data.

The appropriate balance depends on the decision context and consequences of error. In high-impact applications, interpretability may be particularly important because users need to understand and challenge AI recommendations. Model selection should therefore consider both technical performance and practical decision requirements.

8.2 Reliability Under Changing Conditions

AI decision systems may experience declining performance when the environment changes substantially from the conditions represented in training data. Changes in user behavior, operational processes, economic conditions, or technological infrastructure can alter relationships between variables. Model drift can therefore reduce decision reliability over time.

Monitoring data distributions and prediction errors can provide early indications of degradation. Adaptive retraining can restore performance when appropriately managed. A reliable AI system should consequently include mechanisms for detecting and responding to changing conditions.

8.3 Bias and Fairness

AI systems can reproduce or amplify biases present in historical datasets. If certain groups or conditions are underrepresented, predictions may become systematically less reliable for those cases. Bias can reduce both decision accuracy and perceived legitimacy. Organizations should therefore evaluate training datasets for representativeness and assess model performance across relevant groups.

Bias detection should continue after deployment because changing data can create new disparities. Responsible AI governance requires fairness evaluation to be incorporated into the complete decision lifecycle.

9. Applications of AI-Based Decision Systems

9.1 Business and Management

Organizations can use AI decision systems to support demand forecasting, customer analysis, resource allocation, risk assessment, and operational planning. Machine learning can identify patterns in organizational data and generate predictions for future conditions. AI-supported analytics can reduce the time required to process large quantities of information.

Managers can use these insights to compare alternatives and identify potential risks. However, strategic decisions should consider contextual and organizational factors that may not be fully represented in datasets. AI is therefore most effective when used as decision support rather than unrestricted replacement of managerial judgment.

9.2 Financial Decision-Making

Financial organizations can use AI for risk assessment, fraud detection, credit analysis, forecasting, and transaction monitoring. Machine learning can identify unusual patterns across large volumes of financial information. Predictive models can support early identification of potential risk conditions. Automated analysis can also improve response speed for high-volume transactions.

However, financial AI systems require strong governance because inaccurate or biased decisions can have significant consequences. Explainability, fairness, security, and continuous validation are therefore important components of reliable financial decision systems.

9.3 Healthcare Decision Support

AI can support healthcare professionals by analyzing clinical information, medical images, patient records, and risk indicators. Machine learning and deep learning can assist with classification, prediction, and pattern recognition. AI-based decision support may help identify potential risks and prioritize cases for further professional assessment.

However, healthcare decisions involve complex contextual and ethical considerations that cannot always be represented through historical data. AI should therefore complement professional expertise rather than operate without appropriate clinical oversight. Reliability requires rigorous validation and responsible deployment.

9.4 Industrial Decision-Making

Industrial organizations can use AI to support predictive maintenance, quality control, production planning, resource allocation, and process optimization. Machine learning can identify relationships between operational conditions and production outcomes.

AI systems can provide early warnings when equipment behavior indicates potential failure. Computer vision can support automated quality assessment. These capabilities can improve operational consistency and reduce decision latency. Nevertheless, industrial AI must operate within established safety requirements and should include appropriate human supervision for critical decisions.

10. Challenges and Limitations

10.1 Data Quality

Data quality represents one of the most significant determinants of AI decision performance. Missing values, incorrect measurements, inconsistent formats, and biased samples can reduce predictive accuracy. Historical datasets may also fail to capture rare or emerging conditions. Poor data can therefore lead to systematically unreliable decisions.

Data validation, cleaning, integration, and governance should be implemented before model development. Continuous assessment is also necessary because data quality can change after deployment.

10.2 Algorithmic Bias

Algorithmic bias can occur when training data contain historical inequalities, incomplete representation, or systematic measurement differences. AI models may learn these patterns and reproduce them in subsequent decisions. Such bias can affect both accuracy and fairness.

Organizations should use appropriate evaluation procedures to identify performance differences across relevant data groups. Diverse and representative training data can reduce some sources of bias. Nevertheless, fairness requires continuous monitoring because model behavior may change as new data become available.

10.3 Cybersecurity

AI decision systems may become targets for attacks designed to manipulate input data, model behavior, or system outputs. Data poisoning, adversarial inputs, unauthorized access, and other threats can reduce decision reliability. Security vulnerabilities are particularly important when AI systems are connected to operational infrastructure or sensitive information.

Secure data pipelines, access controls, monitoring, and model-protection mechanisms are therefore necessary. Cybersecurity should be incorporated throughout the AI lifecycle rather than added after deployment.

11. Future Research

Future research should investigate hybrid AI architectures that combine machine learning, deep learning, ensemble methods, symbolic reasoning, and explainable AI. Such approaches may improve both predictive performance and transparency by combining complementary computational capabilities.

Research should also develop uncertainty-aware decision systems capable of communicating confidence levels and identifying cases requiring human review. Longitudinal studies could evaluate how AI decision accuracy and reliability evolve after deployment. Comparative research across industries would further clarify which AI techniques are most effective under different decision requirements.

Another important research direction involves developing adaptive and trustworthy AI decision frameworks. Future studies should examine automated methods for detecting model drift, identifying bias, monitoring fairness, and updating models under controlled conditions. Human-AI collaboration should also receive greater attention, particularly regarding appropriate decision authority, user trust, explanation quality, and accountability. Privacy-preserving and federated learning could support AI decision-making across distributed environments without unnecessary data centralization. These developments can contribute to AI systems that are not only accurate but also reliable, transparent, secure, and responsible.

12. Conclusion

Artificial intelligence techniques provide significant opportunities to improve decision accuracy and reliability across diverse organizational and technological environments. Machine learning, deep learning, ensemble methods, predictive analytics, natural language processing, and explainable AI can enhance the ability of decision systems to process information, identify patterns, estimate outcomes, and support evidence-based decisions. The proposed framework emphasizes that reliable AI decision-making requires integration of data quality, predictive modeling, validation, explainability, human oversight, and continuous feedback.

The analysis demonstrates that accuracy and reliability are related but distinct characteristics of intelligent decision systems. A model may generate consistent outputs without necessarily producing correct decisions, particularly when training data are biased or outdated. Similarly, a highly accurate model may become unreliable when environmental conditions change. Continuous monitoring, model validation, uncertainty assessment, and appropriate adaptation are therefore essential. Human oversight remains particularly important for complex or high-impact decisions where contextual knowledge and ethical considerations cannot be fully automated.

Overall, AI can strengthen decision-making by providing scalable analytical capabilities and supporting timely, consistent, and predictive assessments. However, responsible implementation requires organizations to address data quality, bias, explainability, cybersecurity, model drift, and automation risks. Future AI decision systems should focus not only on maximizing predictive accuracy but also on creating trustworthy and accountable decision processes. A balanced combination of artificial intelligence and human expertise can provide a sustainable foundation for accurate, reliable, and responsible decision-making in increasingly complex environments.

References

1. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32.
2. Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5), 1189–1232.
3. Dietterich, T. G. (2000). Ensemble methods in machine learning. In *Multiple classifier systems* (pp. 1–15). Springer.

4. Dietterich, T. G. (2000). An experimental comparison of three methods for constructing ensembles of decision trees: Bagging, boosting, and randomization. *Machine Learning*, 40, 139–157.
5. Roweis, S. T., & Saul, L. K. (2000). Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500), 2323–2326.
6. Cunningham, P., Carney, J., & Jacob, S. (2000). Stability problems with artificial neural networks and the ensemble solution. *Artificial Intelligence in Medicine*, 20(3), 217–225.
7. Cortés, U., Sánchez-Marrè, M., Ceccaroni, L., R-Roda, I., & Poch, M. (2000). Artificial intelligence and environmental decision support systems. *Applied Intelligence*, 13, 77–91.
8. Breese, J. S., Heckerman, D., & Kadie, C. (2000). Empirical analysis of predictive algorithms for collaborative filtering. In *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence* (pp. 43–52). Morgan Kaufmann.
9. Forgionne, G. A., Kohli, R., & Jennings, D. (2002). An AHP analysis of quality in AI and DSS journals. *Omega*, 30(3), 171–183.
10. Adomavicius, G., & Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17(6), 734–749.
11. Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
12. Hinton, G. E., Osindero, S., & Teh, Y.-W. (2006). A fast learning algorithm for deep belief nets. *Neural Computation*, 18(7), 1527–1554.
13. Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning: Data mining, inference, and prediction* (2nd ed.). Springer.
14. Kohavi, R., Longbotham, R., Sommerfield, D., & Henne, R. M. (2009). Controlled experiments on the web: Survey and practical guide. *Data Mining and Knowledge Discovery*, 18, 140–181.
15. Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78–87.
16. Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165–1188.
17. Phillips-Wren, G. (2013). Intelligent decision support systems. In *Multicriteria decision aid and artificial intelligence*. Wiley.
18. LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, 521(7553), 436–444.
19. Ribeiro, M. T., Singh, S., & Guestrin, C. (2016). “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1135–1144). ACM.
20. Sagi, O., & Rokach, L. (2018). Ensemble learning: A survey. *WIREs Data Mining and Knowledge Discovery*, 8(4), e1249.