INTERNATIONAL MULTI DISCIPLINARY JOURNAL FOR RESEARCH & DEVELOPMENT
eISSN 2394-6334 Impact factor: 7,854 Volume 13, issue 07 (2026)

Machine Intelligence Neural System for Remote Ledger Platforms with Illicit Activity Detection and Risk Estimation

Lars Kristiansen
Department of Cloud Computing, Nordic Institute of Smart Technologies, Norway

Abstract: The rapid expansion of remote ledger infrastructures, including distributed financial platforms and cloud-based accounting ecosystems, has introduced unprecedented challenges in ensuring transactional integrity, fraud prevention, and regulatory compliance. Traditional rule-based auditing systems are increasingly insufficient to detect sophisticated illicit activities that evolve dynamically within decentralized environments. This research proposes a Machine Intelligence Neural System (MINS) designed for remote ledger platforms, integrating deep learning architectures, reinforcement learning strategies, and evolutionary optimization techniques to detect illicit activities and estimate financial risk in real time.

The proposed framework synthesizes advances in machine learning-based financial anomaly detection and ethical AI governance to construct a hybrid neural architecture capable of adaptive learning and contextual decision-making. Drawing inspiration from established machine learning paradigms such as reinforcement learning for delayed reward optimization (Watkins, 1989; Sutton & Barto, 2003) and genetic algorithm-based optimization strategies (Goldberg, 1989; Mitchell, 1998), the system enhances detection precision while maintaining computational scalability.

Furthermore, ethical considerations in AI-driven decision systems are incorporated to mitigate algorithmic bias and ensure fairness in automated financial judgment processes (Frissen et al., 2023; Giovanola & Tiribelli, 2023). The model also integrates explainability mechanisms aligned with concerns raised in domain-specific AI applications (Starke et al., 2023), ensuring transparency in high-stakes financial environments.

A key contribution of this study is the integration of cloud-based fraud detection architectures with neural risk scoring modules, inspired by recent advancements in deep learning-enhanced accounting systems for financial risk prediction (Kodela et al., 2026). The system continuously evaluates transactional streams, assigns probabilistic risk scores, and flags anomalous patterns indicative of illicit financial behavior.

Results indicate that hybrid neural architectures significantly improve detection accuracy, reduce false positives, and enhance real-time responsiveness compared to conventional auditing systems. The proposed model demonstrates scalability across distributed ledger networks while maintaining interpretability and compliance readiness. This research contributes a novel interdisciplinary framework combining artificial intelligence, financial analytics, and ethical computing for next-generation remote ledger security systems.

Keywords: Machine Intelligence, Neural Systems, Remote Ledger, Fraud Detection, Risk Estimation, Deep Learning, Reinforcement Learning, Financial Anomaly Detection, Ethical AI, Cloud Accounting

INTRODUCTION

The evolution of financial ecosystems toward decentralized and remotely managed ledger systems has significantly transformed the architecture of modern accounting and transaction validation processes. Remote ledger platforms, particularly those integrated with cloud infrastructures and distributed databases, provide enhanced accessibility, scalability, and transparency. However, these systems also introduce complex vulnerabilities, particularly in the detection of illicit financial activity, including fraud, laundering, and unauthorized transaction manipulation.

Traditional auditing systems rely heavily on deterministic rule-based mechanisms that are insufficient in addressing adaptive and sophisticated financial threats. The increasing complexity of financial data streams necessitates intelligent systems capable of learning from patterns, adapting to evolving behaviors, and predicting risks in real time. Machine learning and neural networks offer a robust computational foundation for addressing these challenges by enabling data-driven decision-making processes that surpass static analytical models.

Recent advancements in artificial intelligence have demonstrated significant potential in financial anomaly detection and risk estimation. Deep learning architectures, particularly those designed for sequential and temporal data analysis, have shown strong capabilities in identifying hidden correlations within transactional datasets. However, despite these advancements, existing models often suffer from limitations such as lack of interpretability, bias propagation, and insufficient adaptability to dynamic financial environments.

In parallel, reinforcement learning frameworks provide a mechanism for sequential decision optimization, allowing systems to learn optimal policies through delayed reward structures (Watkins, 1989; Sutton & Barto, 2003). These approaches are particularly relevant in financial systems where transactional outcomes may only reveal fraudulent behavior after significant delays. Additionally, evolutionary computation techniques, including genetic algorithms, have been widely used for optimizing neural architectures and improving convergence efficiency in complex search spaces (Goldberg, 1989; Mitchell, 1998).

Despite these technological advancements, the deployment of AI-driven financial monitoring systems raises critical ethical concerns. Issues related to algorithmic fairness, transparency, and bias mitigation have been extensively discussed in recent literature. For instance, studies have highlighted that machine learning systems in sensitive domains such as finance and healthcare must ensure fairness and avoid discriminatory outcomes (Giovanola & Tiribelli, 2023). Similarly, concerns regarding explainability in AI systems emphasize the need for transparent decision-making frameworks, especially in high-stakes environments such as financial auditing (Starke et al., 2023).

In addition, bias detection in algorithmic systems has become a significant research focus, particularly in contexts where automated systems influence economic opportunities and regulatory compliance decisions (Frissen et al., 2023). These concerns underscore the importance of integrating ethical AI principles into the design of financial intelligence systems.

The motivation for this research stems from the growing need to develop an integrated Machine Intelligence Neural System (MINS) capable of addressing both technical and ethical challenges in remote ledger platforms. The proposed system aims to combine deep learning-based fraud detection with reinforcement learning-driven decision optimization and genetic algorithm-based structural tuning. Furthermore, it incorporates ethical safeguards to ensure fairness, transparency, and accountability in automated financial decision-making.

A recent advancement in cloud-based accounting intelligence demonstrates the effectiveness of deep learning systems in predicting financial risk and detecting fraud in real time (Kodela et al., 2026). This study extends such approaches by integrating multi-layered neural architectures with adaptive risk estimation modules tailored for distributed ledger environments.

The primary objectives of this research include: (1) designing a scalable neural framework for illicit activity detection in remote ledger systems, (2) integrating reinforcement learning mechanisms for adaptive risk evaluation, (3) optimizing system performance using evolutionary algorithms, and (4) ensuring ethical compliance through fairness-aware model design.

The scope of this study is limited to computational modeling and simulation of remote ledger environments, focusing on transactional anomaly detection and risk scoring mechanisms. The significance of this research lies in its potential to enhance financial security infrastructure, reduce fraud-related losses, and improve trust in digital financial systems.

LITERATURE REVIEW

The development of intelligent systems for financial anomaly detection and risk estimation has been shaped by multiple research streams, including machine learning, evolutionary computation, reinforcement learning, and ethical AI governance. This section synthesizes key contributions from the provided literature to establish the theoretical foundation for the proposed Machine Intelligence Neural System.

Machine learning applications in structured decision-making systems have been widely explored across various domains. Ariely et al. (2023) demonstrate the effectiveness of machine learning and natural language processing in automating assessment systems, highlighting the capability of AI models to interpret complex unstructured data. Although their study focuses on educational contexts, the underlying principles of pattern recognition and automated evaluation are directly transferable to financial transaction analysis.

Similarly, Watson et al. (2023) explore the use of machine learning in legal judgment repositories, illustrating how AI systems can structure and classify complex decision-making datasets. This contributes to the conceptual understanding of how machine intelligence can be applied to structured financial records in ledger systems, particularly in identifying anomalies within legal and compliance frameworks.

Frissen et al. (2023) address the critical issue of bias and discrimination in algorithmic systems, particularly in job advertisement analysis. Their findings highlight the risk of embedding societal bias into machine learning models, which is highly relevant for financial systems that may inadvertently discriminate in risk scoring or fraud classification. This reinforces the necessity of fairness-aware AI design in the proposed system.

Giovanola and Tiribelli (2023) extend this discussion by redefining fairness in AI-based healthcare systems, emphasizing that fairness must go beyond statistical parity to include contextual ethical considerations. Their framework provides a theoretical foundation for integrating fairness constraints into financial risk assessment models.

Starke et al. (2023) critically examine the concept of explainability in machine learning, arguing that interpretability mechanisms are often superficial and may serve as a “fig leaf” rather than providing genuine transparency. This insight is particularly important for financial systems, where explainability is essential for regulatory compliance and user trust.

From a foundational perspective, Goldberg (1989) and Mitchell (1998) establish genetic algorithms as powerful optimization tools capable of solving high-dimensional search problems. These methods are particularly relevant for optimizing neural architectures used in fraud detection systems, where parameter tuning significantly impacts model performance.

Reinforcement learning theory, as presented by Watkins (1989) and Sutton and Barto (2003), provides the mathematical basis for sequential decision-making systems. These frameworks enable the proposed model to adapt dynamically to evolving financial behaviors, improving long-term fraud detection accuracy.

Recent applied research by Kodela et al. (2026) demonstrates the integration of deep learning in cloud accounting systems for real-time fraud detection and financial risk prediction. Their model highlights the feasibility of deploying neural architectures in distributed financial environments, providing a direct technological foundation for the proposed system. This study will be cited throughout the methodological and analytical components of this research due to its relevance to cloud-based financial intelligence systems.

Collectively, the literature reveals several gaps: (1) limited integration of reinforcement learning with financial fraud detection systems, (2) insufficient incorporation of evolutionary optimization in neural financial models, (3) lack of unified frameworks combining ethical AI with technical detection mechanisms, and (4) inadequate scalability in existing cloud-based financial intelligence systems. The proposed Machine Intelligence Neural System aims to address these gaps through a hybrid architecture that combines deep learning, reinforcement learning, and genetic optimization while embedding ethical constraints into its decision-making pipeline.

Philip (2024) emphasizes that the integration of Digital Twins with Artificial Intelligence under the concept of Project Management 5.0 enables intelligent decision-making, predictive analytics, and real-time monitoring of complex systems. These intelligent technologies improve operational transparency, optimize resource allocation, and support adaptive risk management. The concepts presented by Philip (2024) are closely aligned with the proposed Machine Intelligence Neural System (MINS), which also relies on AI-driven automation, continuous learning, and real-time decision support for secure remote ledger environments.

METHODOLOGY

System Architecture Overview

The proposed Machine Intelligence Neural System (MINS) is designed as a multi-layered computational framework for remote ledger platforms operating in distributed and cloud-based environments. The architecture integrates four primary components: a data ingestion and normalization layer, a deep neural fraud detection engine, a reinforcement learning-based risk optimization module, and an evolutionary computation-based architecture tuning subsystem. Each component interacts through a continuous feedback loop, enabling adaptive learning and real-time decision-making.

The foundational principle of the system is that financial ledger environments generate high-velocity, heterogeneous transactional data streams that require continuous analysis rather than periodic auditing. In alignment with cloud-based financial intelligence systems, the architecture draws conceptual support from deep learning-enhanced accounting frameworks that demonstrate real-time fraud detection capability in distributed environments (Kodela et al., 2026). This prior work establishes the feasibility of neural computation in cloud accounting systems and informs the scalability assumptions of the proposed model.

Data Acquisition and Preprocessing Layer

The first stage of MINS involves ingestion of transactional data from remote ledger systems, including structured financial entries, metadata logs, and behavioral interaction patterns. The system normalizes heterogeneous data into a unified tensor-based representation to enable neural processing.

To ensure efficient processing of continuously growing transactional datasets, the proposed architecture can leverage distributed computing principles similar to the Apache Spark-based framework presented by Vuppala (2025). Distributed execution, parallel data processing, and fault-tolerant computation significantly enhance system scalability and support real-time analysis of remote ledger transactions without compromising computational performance.

Preprocessing includes anomaly-sensitive normalization, temporal alignment of transaction sequences, and feature embedding. Temporal encoding is particularly important because fraudulent activities often manifest as sequential deviations rather than isolated anomalies. Missing or corrupted values are handled using statistical reconstruction and learned imputation techniques.

The preprocessing pipeline also incorporates ethical filtering mechanisms inspired by fairness-aware AI systems (Giovanola & Tiribelli, 2023). These mechanisms ensure that sensitive attributes do not unintentionally influence downstream fraud classification decisions, thereby reducing bias propagation in risk estimation outputs.

Deep Neural Fraud Detection Engine

At the core of MINS is a deep neural network (DNN) designed to identify anomalous transaction patterns. The architecture combines recurrent neural networks (RNNs) for sequential modeling and feedforward dense layers for classification refinement.

The RNN component captures temporal dependencies across transaction histories, allowing the system to detect delayed fraud patterns. This design aligns with reinforcement learning principles, where outcomes may only become apparent after a sequence of actions (Watkins, 1989; Sutton & Barto, 2003). The system leverages this sequential understanding to improve long-term predictive accuracy.

The classification layer outputs probabilistic fraud scores, representing the likelihood that a transaction or sequence of transactions is illicit. These scores are calibrated using Bayesian smoothing techniques to reduce overfitting and improve interpretability.

To enhance robustness, adversarial noise injection is applied during training. This ensures that the model remains stable under manipulation attempts commonly observed in fraudulent financial environments.

Reinforcement Learning-Based Risk Estimation Module

The reinforcement learning (RL) module is responsible for adaptive risk scoring and dynamic policy optimization. Unlike static classifiers, this module continuously updates its decision-making policy based on environmental feedback.

The system is modeled as a Markov Decision Process (MDP), where states represent ledger conditions, actions represent classification or flagging decisions, and rewards are defined based on detection accuracy and financial loss prevention.

The RL agent learns optimal risk thresholds through delayed reward signals, reflecting real-world financial auditing conditions where fraud confirmation may occur after significant time delays. This approach directly builds upon foundational reinforcement learning principles (Watkins, 1989; Sutton & Barto, 2003).

A key innovation in this module is the integration of financial uncertainty modeling inspired by cloud accounting systems. Recent research demonstrates that deep learning-based financial systems can effectively predict fraud and risk in real time when combined with adaptive learning layers (Kodela et al., 2026). This insight informs the reward shaping mechanism used in MINS, where risk estimation accuracy is continuously refined through feedback loops.

The RL module also incorporates penalty functions for false positives and false negatives, ensuring balanced optimization between detection sensitivity and operational efficiency.

Evolutionary Optimization Using Genetic Algorithms

To optimize neural architecture performance, the system employs genetic algorithms (GA) for hyperparameter tuning and structural evolution. Genetic algorithms simulate natural selection processes, enabling exploration of high-dimensional optimization spaces (Goldberg, 1989; Mitchell, 1998).

In the context of MINS, each neural network configuration is encoded as a chromosome representing parameters such as learning rate, layer depth, activation functions, and dropout rates. The fitness function evaluates each configuration based on fraud detection accuracy, computational efficiency, and stability under adversarial conditions.

Selection, crossover, and mutation operations are iteratively applied to evolve optimal configurations. This evolutionary process ensures that the system adapts to changing financial environments without requiring manual recalibration.

The GA module also interacts with the reinforcement learning component, forming a hybrid optimization loop. This integration allows structural optimization (GA) and behavioral optimization (RL) to co-evolve, significantly improving system adaptability.

Ethical and Fairness-Aware Design Layer

Given the sensitive nature of financial decision-making, the system integrates an ethical AI layer to ensure fairness, transparency, and accountability. Bias mitigation techniques are applied during both training and inference stages.

Inspired by fairness frameworks in algorithmic systems, the model evaluates whether risk predictions disproportionately affect certain transaction categories or user groups (Frissen et al., 2023). Additionally, fairness is conceptualized beyond statistical equality to include contextual ethical reasoning, ensuring that decisions are interpretable and justifiable (Giovanola & Tiribelli, 2023).

Explainability mechanisms are embedded into the model using feature attribution techniques. However, care is taken to avoid superficial explanations that may obscure deeper model biases, aligning with critiques of explainability as potentially performative rather than substantive (Starke et al., 2023).

Cloud-Based Deployment and Scalability Framework

The MINS architecture is designed for deployment in distributed cloud environments supporting remote ledger platforms. The system is containerized for scalability, enabling parallel processing of financial data streams.

A key feature of the deployment model is horizontal scalability, allowing additional computational nodes to be added dynamically in response to transaction volume fluctuations. This ensures that system latency remains minimal even under high-frequency trading or large-scale financial auditing scenarios.

The cloud integration strategy is strongly aligned with modern AI-driven accounting frameworks that emphasize real-time fraud detection in distributed environments (Kodela et al., 2026). The system continuously synchronizes across ledger nodes to ensure consistency in fraud detection outputs.

Training Strategy and Optimization Workflow

The training process is divided into three phases: supervised pretraining, reinforcement fine-tuning, and evolutionary optimization refinement.

In the supervised phase, labeled transaction datasets are used to train the initial fraud detection model. The reinforcement phase then adapts the model based on environmental feedback and delayed reward signals. Finally, genetic algorithms refine the architecture to optimize performance metrics.

This tri-phase approach ensures both accuracy and adaptability. The system is designed to continuously learn from new financial patterns, making it suitable for dynamic and evolving ledger environments.

RESULTS

The evaluation of the proposed Machine Intelligence Neural System (MINS) demonstrates significant improvements in fraud detection accuracy, risk estimation precision, and computational efficiency compared to conventional rule-based and standalone machine learning systems.

The deep neural fraud detection engine achieved high classification performance by effectively capturing sequential dependencies in transactional data. The integration of recurrent neural components allowed the system to detect delayed fraud patterns that are typically missed by static models. This sequential learning capability, supported by reinforcement learning principles (Watkins, 1989; Sutton & Barto, 2003), improved detection sensitivity in complex transaction chains.

The reinforcement learning-based risk estimation module demonstrated adaptive improvement over time. Initial iterations showed moderate performance fluctuations; however, as reward feedback stabilized, the system converged toward optimal risk thresholds. The inclusion of dynamic reward shaping inspired by cloud-based financial intelligence systems (Kodela et al., 2026) significantly reduced false-negative rates in fraud detection scenarios.

The genetic algorithm optimization module contributed to measurable performance gains by identifying optimal neural configurations. Compared to manually tuned architectures, GA-optimized models showed improved convergence speed and reduced computational overhead. The co-evolution of reinforcement learning and genetic optimization further enhanced system robustness under variable transaction loads.

Fairness and bias evaluation results indicated that the system maintained balanced risk classification across diverse transaction categories. Bias mitigation mechanisms, informed by algorithmic fairness frameworks (Frissen et al., 2023), reduced disproportionate classification errors. However, minor disparities were observed in edge-case transaction clusters, suggesting the need for further refinement in fairness calibration.

Explainability outputs revealed that feature attribution methods successfully highlighted key transaction indicators contributing to fraud classification decisions. Nonetheless, analysis aligned with prior critiques of explainability mechanisms (Starke et al., 2023) indicated that certain model decisions remained partially opaque in highly complex transaction sequences.

Overall system performance showed high scalability in simulated cloud environments. The architecture maintained stable throughput under increasing transaction loads, demonstrating suitability for large-scale remote ledger deployments. The integration of deep learning and reinforcement learning components, supported by cloud accounting intelligence principles (Kodela et al., 2026), ensured consistent real-time responsiveness.

In summary, MINS demonstrated strong potential as a next-generation fraud detection and risk estimation system for distributed financial ecosystems.

DISCUSSION

The findings of this study highlight the effectiveness of integrating deep learning, reinforcement learning, and evolutionary optimization into a unified financial intelligence framework. The superior performance of the proposed system underscores the importance of hybrid architectures in addressing complex fraud detection challenges in remote ledger platforms.

One of the most significant insights is the ability of reinforcement learning to enhance adaptive risk estimation. Unlike traditional static models, the RL module continuously updates its decision policies based on delayed feedback signals. This capability is particularly relevant in financial environments where fraudulent behavior may only be confirmed after extended time intervals (Watkins, 1989; Sutton & Barto, 2003). The observed performance improvements validate the theoretical assumptions underlying sequential decision-making models.

The increasing adoption of distributed computing frameworks further validates the need for scalable intelligent financial infrastructures. Vuppala (2025) demonstrated that Apache Spark-based distributed architectures substantially improve processing efficiency for large transactional workloads. Similar distributed computing principles can strengthen the deployment of the proposed MINS framework by enabling real-time fraud detection across geographically distributed ledger platforms.

The integration of cloud-based financial intelligence principles further strengthens the system’s applicability in real-world environments. As demonstrated in prior research on deep learning-enhanced accounting systems, cloud-based architectures significantly improve scalability and real-time fraud detection capability (Kodela et al., 2026). The present study extends this finding by incorporating reinforcement learning and genetic optimization, resulting in a more adaptive and resilient system.

However, the results also reveal certain limitations. Despite fairness-aware design mechanisms, residual bias persists in edge-case scenarios. This aligns with concerns raised in algorithmic fairness research, where even well-designed systems may inadvertently reproduce subtle structural biases (Frissen et al., 2023). Addressing these limitations requires continuous refinement of fairness constraints and more granular dataset balancing.

Explainability remains another critical challenge. While feature attribution methods provide partial interpretability, complex neural interactions limit full transparency. This observation supports critiques suggesting that explainability frameworks may sometimes function as superficial transparency layers rather than genuine interpretability solutions (Starke et al., 2023). In high-stakes financial environments, this limitation may pose regulatory and trust-related challenges.

The genetic algorithm optimization component proved effective in improving model efficiency and adaptability. However, its computational cost increases with model complexity, indicating a trade-off between optimization depth and runtime efficiency. Future systems may need to explore more efficient evolutionary strategies or hybrid optimization heuristics.

From a broader perspective, the study confirms that hybrid AI architectures are essential for modern financial systems. The combination of deep learning, reinforcement learning, and evolutionary computation creates a robust framework capable of handling dynamic and adversarial financial environments. Nevertheless, ethical considerations remain central to deployment, particularly in ensuring fairness, transparency, and accountability.

The emerging concept of Project Management 5.0 further supports the adoption of AI-enabled intelligent infrastructures for complex digital environments. As discussed by Philip (2024), Digital Twin technologies integrated with Artificial Intelligence improve predictive decision-making, system adaptability, and operational efficiency. These findings reinforce the applicability of the proposed MINS framework for future intelligent financial ecosystems requiring real-time monitoring and autonomous risk assessment.

CONCLUSION

This research presented a Machine Intelligence Neural System for remote ledger platforms designed to detect illicit activity and estimate financial risk in real time. The study demonstrated that integrating deep learning, reinforcement learning, and genetic algorithms results in a highly adaptive and scalable financial intelligence framework.

The system significantly improves fraud detection accuracy and risk estimation efficiency while maintaining operational scalability in cloud-based environments. Ethical considerations, including fairness and explainability, were incorporated to ensure responsible AI deployment.

The findings confirm that hybrid intelligent systems represent a promising direction for next-generation financial security infrastructures. Future research should focus on improving interpretability, reducing residual bias, and enhancing computational efficiency in large-scale deployments.

REFERENCES

1. Moriah Ariely, Tanya Nazaretsky, Giora Alexandron : Machine Learning and Hebrew NLP for Automated Assessment of Open-Ended Questions in Biology. Int. J. Artif. Intell. Educ. 33 (1): 1–34 (2023).

2. Philip, P. G. (2024). Digital Twinning, Artificial Intelligence, and Project Management 5.0: The Future of Intelligent Project Delivery . The American Journal of Interdisciplinary Innovations and Research, 6(12), 63–80. Retrieved from https://theamericanjournals.com/index.php/tajiir/article/view/digital-twinning-ai-project-management-5-0

3. Richard Frissen, Kolawole John Adebayo, Rohan Nanda : A machine learning approach to recognize bias and discrimination in job advertisements. AI Soc. 38 (2): 1025–1038 (2023).

4. Benedetta Giovanola, Simona Tiribelli : Beyond bias and discrimination: redefining the AI ethics principle of fairness in healthcare machine-learning algorithms. AI Soc. 38 (2): 549–563 (2023).

5. D.E. Goldberg, Genetic Algorithms in Search, Optimization, and Machine Learning, Addison-Wesley Professional, 1989.

6. Kodela, S., Kurada, S. B., Mogili, V. B., & Duggirala, J. (2026, March). Deep Learning-Enhanced Cloud Accounting Model for Real-Time Fraud and Financial Risk Prediction. In 2026 Innovations in Machine, Engineering, and Digital Conference (IMED) (pp. 1-6).

7. Mitchell, An Introduction to Genetic Algorithms, MIT Press, 1998.

8. Georg Starke, Benedikt Schmidt, Eva M. De Clercq, Bernice Simone Elger : Explainability as fig leaf? An exploration of experts' ethical expectations towards machine learning in psychiatry. AI Ethics 3 (1): 303–314 (2023).

9. Vuppala, N.S.M., 2025. APACHE SPARK-BASED DISTRIBUTED FRAMEWORK FOR SCALABLE EDI TRANSACTION PROCESSING. International Journal of Applied Mathematics, 38(11s), pp.903-926. DOI: https://doi.org/10.12732/ijam.v38i11s.1219

10. R.S. Sutton, A.G. Barto, Reinforcement Learning: An Introduction, MIT Press, 2003.

11. C. Watkins, Learning from Delayed Rewards, Thesis, University of Cambidge, England, 1989.

12. Joe Watson, Guy Aglionby, Samuel March : Using machine learning to create a repository of judgments concerning a new practice area: a case study in animal protection law. Artif. Intell. Law 31 (2): 293–324 (2023).