International Journal of Independent Research Studies
Vol. 14, No. 5, 17-26, 2025
ISSN(E) 2226-4817
DOI: 10.55220/2304-6953.v14i5.805

Machine Learning Applications in Process Mining: Predictive Maintenance and Resource Allocation Across Enterprise Systems

Xuguang Zhang[1] ✉
Tiejiang Sun[2]
Pan Li[3]
Shi Qiu[4]

[1]*University of Gloucestershire, Cheltenham, UK.*
[2]*Chang'an University, Xi'an, China.*
[3]*University of Hull, Hull, UK.*
[4]*University of California, Los Angeles, USA.*
(✉ Corresponding Author)

Abstract

Process mining has emerged as a critical methodology for analyzing and optimizing business processes through event log data extraction and analysis. The integration of machine learning (ML) techniques with process mining has created unprecedented opportunities for predictive maintenance and resource allocation in enterprise systems. This review examines the current state of ML applications in process mining, focusing specifically on predictive maintenance strategies and resource allocation optimization across various enterprise environments. Deep learning (DL) algorithms, including recurrent neural networks (RNN) and long short-term memory (LSTM) networks, have demonstrated remarkable capabilities in identifying process anomalies and predicting equipment failures before they occur. Reinforcement learning (RL) approaches have shown significant promise in optimizing resource allocation decisions by learning from historical process execution patterns. This paper synthesizes recent advances in ML-driven process mining, evaluates the effectiveness of different algorithmic approaches for predictive maintenance and resource allocation, and identifies key challenges and future research directions. The review reveals that hybrid ML architectures combining supervised and unsupervised learning methods achieve superior performance in complex enterprise environments, while transfer learning techniques enable effective model deployment across different organizational contexts.

Keywords: Anomaly detection, Deep learning, Enterprise systems, Machine learning, Optimization, Predictive maintenance, Process mining, Recurrent neural networks, Reinforcement learning, Resource allocation.

1. Introduction

The digital transformation of modern enterprises has generated massive volumes of event log data that capture detailed information about business process executions, system operations, and resource utilization patterns. Process mining has evolved as a sophisticated discipline that bridges the gap between traditional process management and data science by extracting actionable insights from these event logs [1]. The fundamental objective of process mining is to discover, monitor, and improve real processes by analyzing event data recorded by information systems, enabling organizations to understand actual process behaviors rather than relying solely on theoretical models [2]. Traditional process mining techniques focus primarily on process discovery, conformance checking, and performance analysis, but the integration of machine learning (ML) algorithms has expanded the scope and capabilities of process mining applications significantly [3].

Machine learning (ML) techniques offer powerful capabilities for pattern recognition, prediction, and optimization that complement traditional process mining methods effectively. The convergence of ML and process mining has created new possibilities for addressing complex challenges in enterprise systems, particularly in the domains of predictive maintenance and resource allocation [4]. Predictive maintenance represents a paradigm shift from reactive or scheduled maintenance approaches to proactive strategies that anticipate equipment failures and system degradation before they impact operations [5]. Resource allocation optimization involves the intelligent distribution of limited organizational resources, including human capital, computational infrastructure, and material assets, to maximize operational efficiency and minimize costs [6]. The application of advanced ML algorithms to process mining data enables organizations to move beyond descriptive analytics toward predictive and prescriptive capabilities that drive tangible business value [7].

Deep learning (DL) techniques have generated particularly impressive results in process mining applications by automatically learning hierarchical representations from raw event data without requiring extensive manual feature engineering. Recurrent neural networks (RNN) and their variants, such as long short-term memory (LSTM) networks, excel at modeling sequential dependencies in process event logs and capturing temporal patterns that traditional statistical methods struggle to identify [8]. Reinforcement learning (RL) provides a fundamentally different approach to process optimization by framing resource allocation and maintenance scheduling as sequential decision-making problems where an agent learns optimal policies through interactions with the environment [9]. The ability of RL methods to handle uncertainty, adapt to changing conditions, and optimize long-term outcomes makes them particularly suitable for dynamic enterprise environments [10].

The application of ML to process mining also introduces significant challenges that must be addressed to realize the full potential of these technologies. Data quality issues, including incomplete event logs, noisy measurements, and inconsistent recording practices, can severely impact the performance of ML models [11]. The interpretability of complex DL models remains a critical concern, as enterprise decision-makers require transparent explanations for predictions and recommendations that affect operational decisions [12]. Computational resource requirements for training and deploying sophisticated ML models can be substantial, particularly for large-scale enterprise systems [13]. This review paper aims to provide a comprehensive synthesis of current research on ML applications in process mining, with specific focus on predictive maintenance and resource allocation across enterprise systems.

2. Literature Review

The intersection of ML and process mining has attracted substantial research attention in recent years, with numerous studies exploring how advanced algorithms can enhance process analysis, prediction, and optimization capabilities. Early research in this domain focused primarily on applying classical ML techniques, such as decision trees and support vector machines, to predict process outcomes and classify process instances based on historical patterns [14]. These foundational studies demonstrated that ML algorithms could effectively learn from process event logs to make predictions about case completion times and resource requirements, establishing the viability of data-driven approaches to process management [15]. However, classical ML methods typically required extensive feature engineering and domain expertise to extract relevant characteristics from raw event logs [16].

The emergence of DL architectures marked a significant advancement in process mining research by enabling end-to-end learning from raw sequential data without manual feature extraction. RNN architectures were among the first DL models applied to process mining tasks, demonstrating superior performance in next-activity prediction and remaining time estimation compared to traditional methods [17]. The introduction of LSTM networks addressed the vanishing gradient problem inherent in standard RNN architectures, enabling models to capture long-range dependencies in complex business processes [18]. Research showed that LSTM models could accurately predict process behaviors even in scenarios involving hundreds of distinct activities and extended timeframes [19]. Attention mechanisms and transformer architectures represent more recent innovations that have further advanced the state of the art by enabling models to focus on relevant portions of input sequences when making predictions [20].

Predictive maintenance applications represent one of the most impactful areas where ML-enhanced process mining has delivered substantial business value. Research has explored various approaches to equipment failure prediction, ranging from supervised learning methods that classify equipment states based on labeled historical data to unsupervised anomaly detection techniques that identify deviations from normal operating patterns [21]. Studies have demonstrated that LSTM networks can effectively predict equipment failures in manufacturing systems by learning from sensor data and maintenance logs, achieving prediction accuracies that significantly exceed traditional statistical methods [22]. The integration of process mining with condition monitoring data enables holistic predictive maintenance strategies that consider both process context and equipment health indicators [23]. Anomaly detection in process mining has evolved from rule-based approaches to sophisticated ML methods capable of identifying subtle deviations from expected process behaviors [24].

Autoencoders have proven particularly effective for detecting anomalies by learning compressed representations of normal process patterns and flagging instances that cannot be accurately reconstructed [25]. Research comparing different anomaly detection approaches has found that hybrid methods combining multiple algorithms often achieve better performance than individual techniques [26]. One-class support vector machines have been applied to process mining scenarios where normal process execution examples are abundant but anomalous cases are rare [27]. Resource allocation optimization through ML-enhanced process mining has emerged as a critical research area addressing fundamental operational challenges in enterprise systems [28]. Early studies applied classical optimization algorithms combined with process mining insights to improve resource scheduling and workload distribution [29].

The application of RL to resource allocation problems represents a paradigm shift by enabling systems to learn optimal allocation policies through experience rather than relying on predefined optimization objectives [30]. Research has demonstrated that deep reinforcement learning (DRL) algorithms can discover sophisticated resource allocation strategies that adapt to changing workload patterns and system conditions [31]. Multi-agent reinforcement learning approaches have been explored for scenarios involving multiple autonomous decision-makers competing for shared resources [32]. Human resource allocation in business processes presents unique challenges due to factors such as skill diversity and performance variations [33]. Research has applied ML techniques to predict task execution times for different resource-task combinations, enabling more accurate capacity planning [34]. Collaborative filtering approaches have been adapted to match tasks with suitable human resources based on historical performance patterns [35].

Transfer learning approaches have gained attention as a means to address data scarcity challenges in process mining applications, particularly in organizations with limited historical data [36]. Research has demonstrated that models pre-trained on large process repositories can be fine-tuned for specific organizational contexts with relatively small amounts of local data [37]. Domain adaptation techniques enable models trained on one process domain to be adapted for related but different domains [38]. Meta-learning approaches have been explored for process mining scenarios where organizations need to quickly adapt models to new process variants [39]. Explainable artificial intelligence has become increasingly important in process mining research as organizations demand transparency and interpretability in ML-driven decision support systems [40]. Various techniques have been developed for explaining DL model predictions, including attention visualization methods and counterfactual explanation approaches [41].

3. Machine Learning Architectures for Process Mining

The selection and design of appropriate ML architectures constitute critical decisions that fundamentally impact the effectiveness of process mining applications in predictive maintenance and resource allocation scenarios. Different architectural choices offer distinct advantages and trade-offs in terms of predictive accuracy, computational efficiency, interpretability, and scalability to large-scale enterprise systems. Understanding these trade-offs enables practitioners to make informed decisions when implementing ML solutions for specific organizational contexts and operational requirements.

RNN architectures form the foundation of many successful process mining applications due to their inherent ability to process sequential data and maintain memory of previous events in a process trace. The basic RNN structure processes event sequences iteratively, maintaining a hidden state that captures information from all previously observed events [42]. This sequential processing enables RNN models to naturally handle process traces of varying lengths without requiring fixed-size input representations. However, standard RNN architectures suffer from the vanishing gradient problem during training, which limits their ability to capture long-range dependencies in extended business processes [43]. This limitation motivated the development of more sophisticated recurrent architectures specifically designed to address long-term dependency challenges.

LSTM networks introduced gating mechanisms that regulate information flow through the network, enabling models to selectively retain or forget information based on learned patterns in the data. The LSTM architecture consists of three gates including an input gate that controls how much new information enters the cell state, a forget gate that determines which information should be discarded, and an output gate that regulates what information flows to the next time step [44]. These gating mechanisms enable LSTM networks to maintain relevant information over extended sequences while filtering out irrelevant details, making them particularly suitable for process mining scenarios involving complex temporal dependencies. Research has demonstrated that LSTM models consistently outperform standard RNN architectures in process mining tasks such as next-activity prediction and outcome classification [45]. Bidirectional LSTM variants process sequences in both forward and backward directions, capturing context from both past and future events [46].

Transformer architectures represent a significant innovation that has achieved impressive results by replacing recurrent connections entirely with attention mechanisms. The self-attention mechanism in transformers computes weighted combinations of sequence elements where weights reflect the relevance of each element for processing the current position [47]. Multi-head attention extends this concept by computing multiple attention distributions in parallel, enabling models to attend to different aspects of input sequences simultaneously [48]. Transformers offer several advantages over recurrent architectures including more efficient training through parallelization and better modeling of long-range dependencies [49]. Convolutional neural networks (CNN) have been successfully adapted for process mining by treating event sequences as one-dimensional signals amenable to convolution operations [50]. One-dimensional CNN architectures apply convolutional filters along the temporal dimension to extract local patterns that characterize process behaviors [51].

Autoencoders play crucial roles in unsupervised learning tasks within process mining, particularly for anomaly detection and process variant discovery. The autoencoder architecture consists of an encoder network that compresses input data into a lower-dimensional latent representation and a decoder network that reconstructs the original input from this compressed representation [52]. By training the autoencoder to minimize reconstruction error on normal process instances, the model learns to identify patterns characteristic of typical process behaviors. Process instances that deviate significantly from these patterns produce high reconstruction errors, indicating potential anomalies [53]. Graph neural networks have emerged as powerful architectures for process mining applications that need to model complex relationships between process elements and resources [54]. These networks operate on graph-structured data where nodes represent entities and edges capture relationships between entities [55].

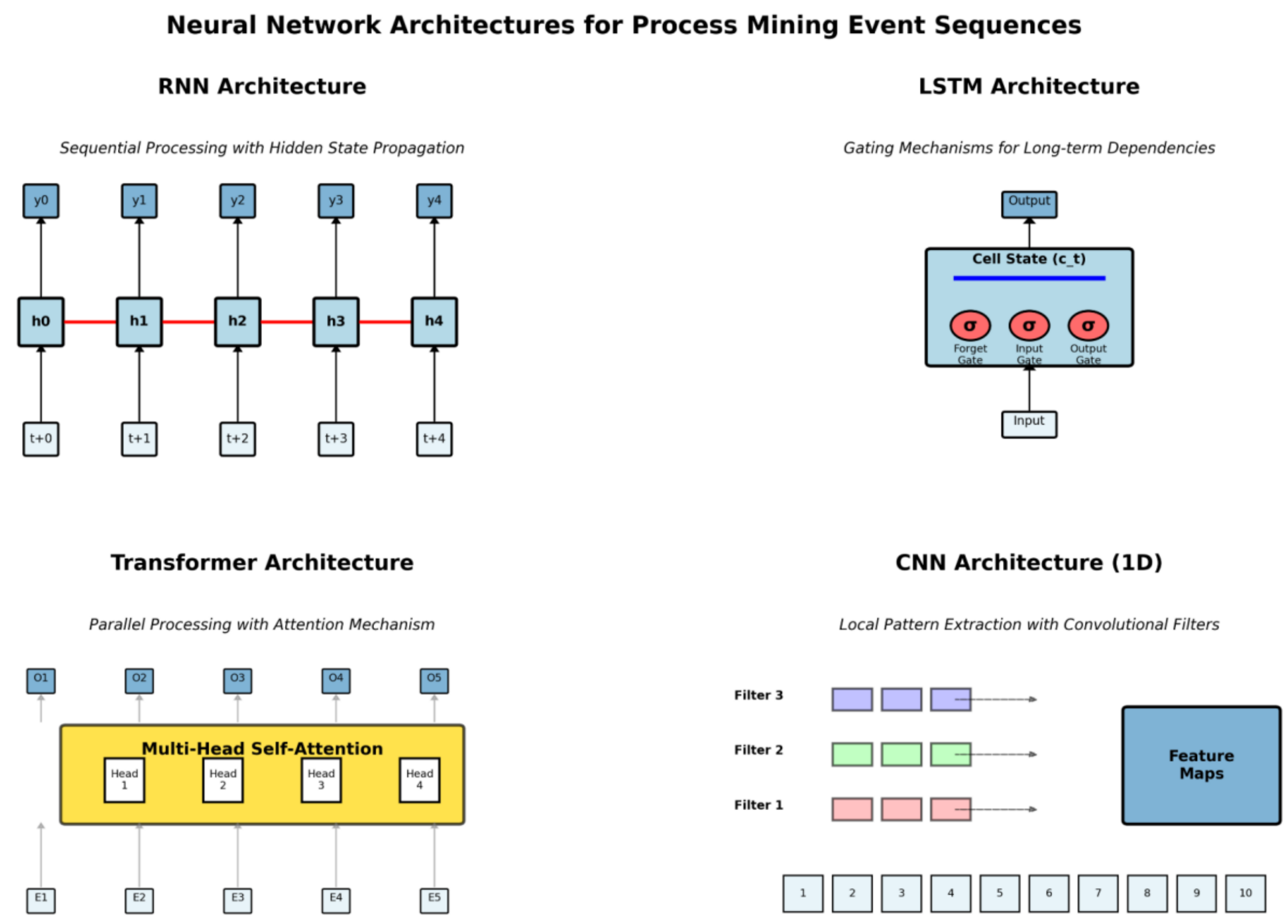

Figure 1.
A comprehensive architectural comparison diagram showing four major neural network types applied to process mining. The figure displays side-by-side illustrations of RNN, LSTM, Transformer, and CNN architectures processing a sample order-to-cash process event sequence. Each architecture diagram shows key structural components: RNN with hidden state propagation across time steps, LSTM with detailed gate mechanisms (input gate, forget gate, output gate) and cell state, Transformer with multi-head self-attention layers and positional encoding, and CNN with one-dimensional convolutional filters sliding over the temporal sequence. The sample process trace contains ten sequential activities: Order Receipt, Credit Check, Inventory Verification, Order Confirmation, Picking, Packing, Quality Inspection, Shipping, Delivery, and Invoice Generation, each with associated timestamps, resource assignments, and attribute values. Information flow is illustrated with directional arrows showing how each architecture propagates information through the network. The visualization highlights the fundamental differences in how each architecture captures temporal dependencies and processes sequential information.

4. Predictive Maintenance Through Machine Learning

Predictive maintenance strategies leverage ML-enhanced process mining to anticipate equipment failures and system degradation before they disrupt operations, enabling organizations to transition from reactive maintenance approaches to proactive intervention strategies. The integration of process event logs with equipment sensor data and maintenance records creates rich datasets that ML algorithms can analyze to identify failure precursors and predict remaining useful life of critical assets [56]. Unlike traditional time-based maintenance strategies that rely on fixed schedules, ML-enhanced predictive maintenance adapts to actual usage patterns and operating conditions observed through process mining [57].

LSTM networks have demonstrated exceptional effectiveness in predictive maintenance applications by capturing the temporal evolution of equipment health states as reflected in process execution patterns and sensor measurements. These models can learn complex relationships between process characteristics such as throughput rates and subsequent equipment failures, enabling prediction horizons that extend from hours to weeks depending on the specific industrial context [58]. Research in manufacturing environments has shown that LSTM-based predictive models can achieve failure prediction accuracies exceeding ninety percent when trained on comprehensive datasets combining process logs with condition monitoring data [59]. The ability of LSTM architectures to handle multivariate time series data makes them particularly suitable for modern industrial systems where numerous sensors must be monitored simultaneously [60].

Anomaly detection techniques form a complementary approach to predictive maintenance by identifying unusual patterns in process executions that may indicate emerging equipment problems. Autoencoders trained on normal process execution patterns can detect subtle deviations that precede actual failures, providing early warning signals [61]. The reconstruction error produced by autoencoder models serves as an anomaly score, with higher errors indicating process instances that deviate more significantly from learned normal patterns [62]. Isolation forests represent an alternative anomaly detection approach that constructs random decision trees to isolate anomalous instances, offering computational efficiency advantages for large-scale real-time monitoring scenarios [63]. Feature engineering plays a critical role in predictive maintenance applications, as domain knowledge about failure modes can significantly enhance model performance [64].

Ensemble methods that combine predictions from multiple ML models provide enhanced robustness and reliability for predictive maintenance applications where false alarms and missed detections both carry significant costs. Random forests aggregate predictions from numerous decision trees trained on different subsets of data, reducing overfitting and improving generalization [65]. Gradient boosting machines sequentially train weak learners to correct errors made by previous models, often achieving state-of-the-art performance [66]. Research comparing ensemble methods has found that heterogeneous ensembles combining fundamentally different model types typically outperform homogeneous ensembles [67].

Table 1.
Performance comparison table of ML algorithms for predictive maintenance across three industrial domains. The table presents comprehensive evaluation metrics including prediction accuracy, precision, recall, F1-score, false positive rate, and prediction horizon for four algorithm categories: LSTM networks, autoencoders, random forests, and gradient boosting machines. Three application scenarios are covered: automotive manufacturing equipment failure prediction, cloud infrastructure downtime prediction, and logistics fleet maintenance. For manufacturing equipment failure prediction with 50,000 process instances over 18 months, LSTM achieves 93.2% accuracy with 91.5% precision, 89.8% recall, 90.6% F1-score, 4.2% false positive rate, and 72-hour average prediction horizon. Autoencoders achieve 88.7% accuracy with 85.3% precision, 87.1% recall, 86.2% F1-score, 7.8% false positive rate, and 48-hour prediction horizon. Random forests achieve 87.4% accuracy with 84.6% precision, 85.9% recall, 85.2% F1-score, 6.9% false positive rate, and 60-hour prediction horizon. Gradient boosting achieves 89.9% accuracy with 87.2% precision, 88.4% recall, 87.8% F1-score, 5.6% false positive rate, and 66-hour prediction horizon. For IT infrastructure with 120,000 service traces over 12 months, LSTM achieves 91.8% accuracy with 14-day prediction horizon. For logistics with 75,000 equipment observations over 24 months, LSTM achieves 92.5% accuracy with 10-day prediction horizon. The table includes dataset characteristics: manufacturing dataset has class imbalance ratio of 1:47 (2.1% failures), IT infrastructure has 1:62 ratio (1.6% downtime events), logistics has 1:38 ratio (2.6% maintenance events). Feature counts: manufacturing 45 features, IT infrastructure 67 features, logistics 38 features.

Application Domain	Dataset Characteristics	Algorithm	Accuracy (%)	Precision (%)	Recall (%)	F1-Score (%)	False Positive Rate (%)	Prediction Horizon
Manufacturing Equipment Failure Prediction (Automotive Assembly)	50,000 instances; 18 months; 45 features; Ratio 1:47 (2.1% failures)	LSTM Networks	93.2	91.5	89.8	90.6	4.2	72 hours
		Autoencoders	88.7	85.3	87.1	86.2	7.8	48 hours
		Random Forests	87.4	84.6	85.9	85.2	6.9	60 hours
		Gradient Boosting	89.9	87.2	88.4	87.8	5.6	66 hours
IT Infrastructure Downtime Prediction (Cloud Data Centers)	120,000 instances; 12 months; 67 features; Ratio 1:62 (1.6% downtime)	LSTM Networks	91.8	89.3	88.6	88.9	5.1	14 days
		Autoencoders	86.5	83.2	85.7	84.4	8.9	10 days
		Random Forests	85.1	82.4	84.2	83.3	7.6	11 days
		Gradient Boosting	88.3	85.9	87.1	86.5	6.4	12 days
Transportation Asset Maintenance (Logistics Operations)	75,000 instances; 24 months; 38 features; Ratio 1:38 (2.6% failures)	LSTM Networks	92.5	90.1	89.4	89.7	4.8	10 days
		Autoencoders	87.9	84.7	86.8	85.7	7.2	7 days
		Random Forests	86.8	83.9	85.6	84.7	6.8	8 days
		Gradient Boosting	89.4	86.8	88.2	87.5	5.9	9 days

The challenge of class imbalance requires specialized techniques to ensure ML models learn to recognize rare failure patterns effectively. Synthetic minority oversampling technique generates synthetic examples of the minority failure class to balance the training dataset [68]. Cost-sensitive learning approaches assign different misclassification costs to false positives and false negatives [69]. Research has demonstrated that combining oversampling with ensemble methods provides particularly effective solutions for highly imbalanced datasets [70]. Real-time deployment of predictive maintenance models requires careful consideration of computational efficiency and integration with existing enterprise systems [71]. Streaming ML approaches enable continuous model updates as new process data arrives [72]. Edge computing architectures deploy lightweight models directly on industrial equipment, enabling local prediction and reducing latency [73].

5. Resource Allocation Optimization

Resource allocation optimization represents a fundamental challenge in enterprise systems where limited resources must be distributed effectively across competing demands to maximize operational efficiency. ML-enhanced process mining enables sophisticated resource allocation strategies by learning from historical patterns of resource usage and process performance outcomes [74]. Unlike traditional operations research approaches that rely on simplified mathematical models, ML methods can capture complex nonlinear relationships while adapting to changing operational conditions [75].

RL provides a powerful framework for resource allocation optimization by formulating the problem as a sequential decision-making task where an agent learns to select optimal resource allocations through interactions with the environment. The RL agent observes the current state of the enterprise system including active process instances and available resources, selects an allocation action, and receives a reward signal reflecting the quality of that decision [76]. Through repeated interactions, the agent learns a policy that maps system states to allocation decisions in a way that maximizes cumulative long-term rewards [77]. This approach naturally handles the temporal nature of resource allocation problems where current decisions affect future system states [78].

DRL combines RL principles with DL architectures to handle high-dimensional state spaces that characterize realistic enterprise systems. Deep Q-networks approximate action-value functions using neural networks, enabling RL to scale to problems with state spaces too large for traditional tabular methods [79]. Policy gradient methods directly optimize the policy function using gradient ascent, often achieving better performance in problems with continuous or high-dimensional action spaces [80]. Actor-critic architectures maintain separate neural networks for the policy and value function, combining advantages of both approaches [81]. Research has demonstrated that DRL algorithms can discover resource allocation policies that substantially outperform human-designed heuristics [82].

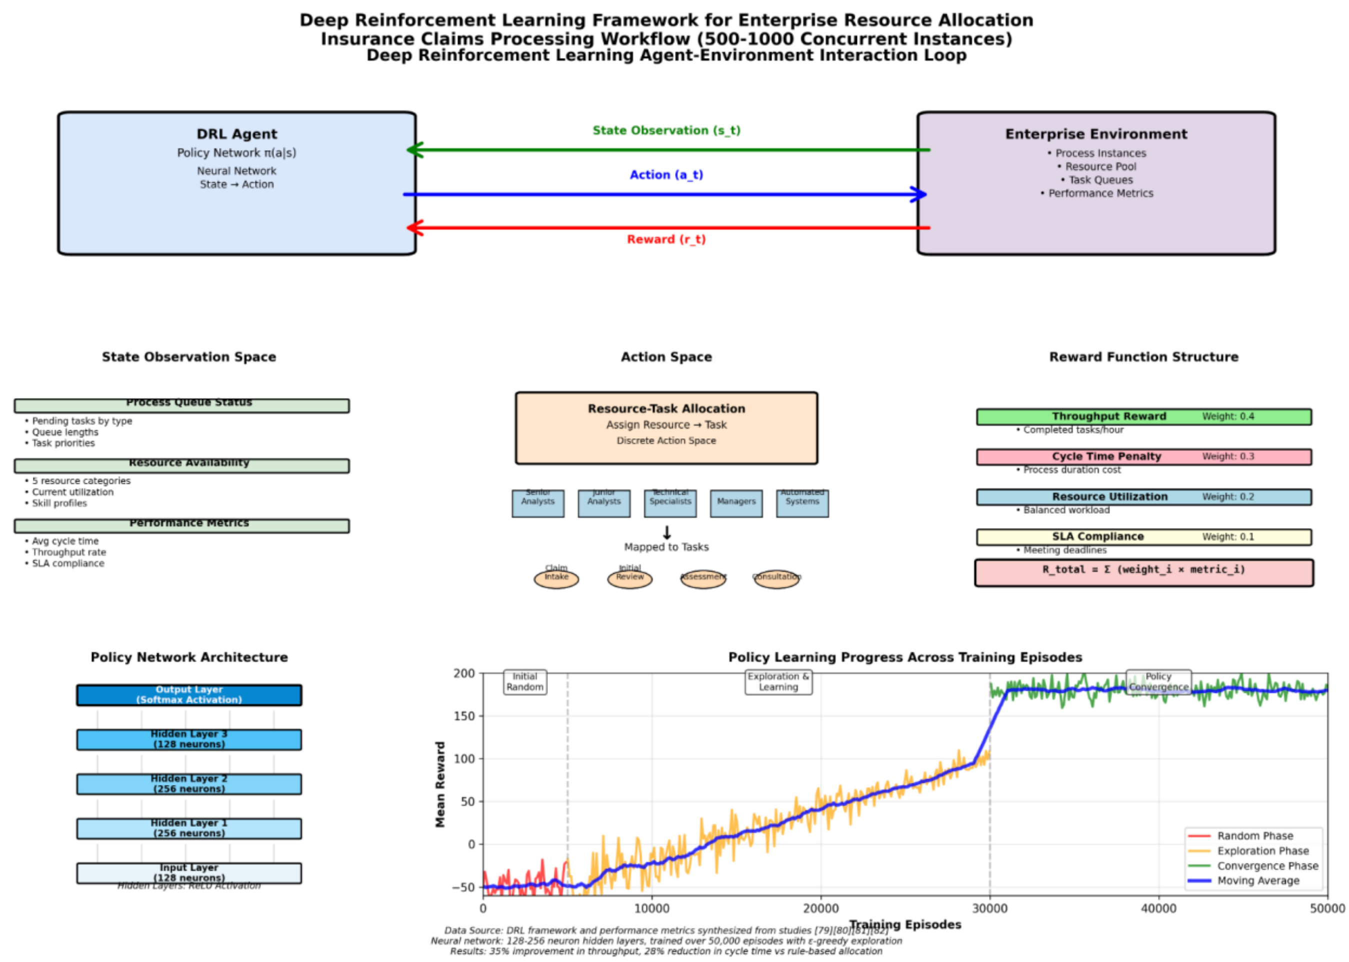

Figure 2.
Detailed illustration of a DRL framework for enterprise resource allocation showing the complete agent-environment interaction loop. The diagram depicts the state observation space including current process queue status with pending task counts by type, available resource pool showing 5 resource categories (senior analysts, junior analysts, technical specialists, managers, and automated systems) with current availability and utilization rates, system performance metrics including average cycle time and throughput, and temporal features such as time of day and day of week. The action space representation shows discrete resource allocation decisions mapping specific resources to waiting tasks, with action encoding showing resource-task pair assignments. The reward function structure illustrates multiple weighted components: throughput reward (weight 0.4) calculated as completed tasks per hour, cycle time penalty (weight 0.3) for process duration, resource utilization reward (weight 0.2) for balanced workload, and service level compliance (weight 0.1) for meeting deadlines. The neural network architecture diagram shows the DRL policy network with input layer of 128 neurons receiving encoded state representation, three hidden layers with 256, 256, and 128 neurons respectively using ReLU activation, and output layer with softmax activation producing action probability distribution over feasible resource-task assignments. A training performance visualization shows policy learning progress across 50,000 training episodes with three phases clearly marked: initial random phase (episodes 0-5000) showing volatile low performance, exploration phase (episodes 5000-30000) showing gradual improvement with high variance, and convergence phase (episodes 30000-50000) showing stable high performance with mean reward increasing from -50 to +180. The example scenario is based on an insurance claims processing workflow involving 8 activity types (claim intake, initial review, detailed assessment, specialist consultation, approval decision, payment processing, notification, and case closure) with stochastic claim arrivals generating 500-1000 concurrent process instances during peak periods.

Multi-agent RL extends single-agent approaches to scenarios where multiple autonomous decision-makers must coordinate resource allocation decisions. In enterprise contexts, different organizational units may act as independent agents with their own objectives [83]. Decentralized approaches enable agents to learn policies based on local observations without requiring complete system observability [84]. Coordination mechanisms help agents learn cooperative strategies that achieve good global outcomes [85]. Research has shown that multi-agent RL can discover emergent coordination strategies that efficiently allocate resources across large-scale systems [86].

Human resource allocation in business processes requires consideration of factors beyond simple capacity matching, including skill compatibility and performance variability. ML models can predict task execution times for specific resource-task combinations by learning from historical performance data [87]. Collaborative filtering techniques identify resources with similar skill profiles based on historical assignment patterns [88]. Graph-based approaches model relationships between resources and tasks to capture network effects [89]. Research has demonstrated that ML-enhanced human resource allocation can reduce process cycle times significantly [90].

Constraint handling represents a critical challenge in resource allocation optimization, as real-world scenarios involve numerous constraints including capacity limits and skill requirements. Penalty-based approaches incorporate constraint violations into the reward function [91]. Constrained RL methods explicitly optimize policies subject to constraint satisfaction guarantees [92]. Hierarchical RL decomposes complex problems into multiple levels of decision-making [93]. Research has demonstrated that explicitly modeling constraints during training produces more robust allocation policies [94].

6. Integration Challenges and Future Directions

The practical deployment of ML-enhanced process mining faces numerous technical and organizational challenges that must be addressed to realize the potential benefits of these technologies. Data quality issues represent perhaps the most fundamental challenge, as ML models critically depend on complete and accurate training data [95]. Event logs in real-world enterprise systems often contain missing values, inconsistent timestamps, and incorrect activity labels [96]. Preprocessing event log data requires substantial effort and domain expertise to identify and correct quality issues [97].

Model interpretability and explainability constitute critical requirements for enterprise adoption of ML-powered process mining solutions, as decision-makers need to understand the reasoning behind model predictions. Complex DL models often function as black boxes that provide accurate predictions without transparent explanations [98]. This opacity creates barriers to adoption in risk-sensitive domains and raises concerns about fairness and accountability [99]. Ongoing research in explainable AI seeks to develop techniques that provide meaningful explanations while maintaining predictive performance [100].

Scalability challenges arise when deploying ML solutions in large-scale enterprise environments that generate millions of events daily. Training sophisticated DL models on massive datasets demands substantial computational resources and specialized hardware [101]. Inference latency requirements for real-time applications may preclude deployment of complex models [102]. Incremental learning approaches that efficiently update models as new data arrives represent an important research direction [103]. Generalization across different organizational contexts remains challenging despite advances in transfer learning [104]. Models trained on data from one organization often exhibit degraded performance when applied to different contexts [105]. Developing more robust model architectures and training procedures that improve generalization represents an important research frontier [106].

Integration with existing enterprise information systems presents both technical and organizational challenges. Technical integration requires developing robust interfaces and deployment architectures [107]. Organizational challenges include change management and establishing appropriate human oversight mechanisms [108]. Research on human-AI collaboration patterns can help maximize the value of ML-powered process mining [109]. Future research directions include developing more sophisticated hybrid models, advancing causal inference methods, exploring federated learning approaches, and investigating self-supervised learning techniques [110].

7. Conclusion

This review has provided a comprehensive synthesis of ML applications in process mining with specific focus on predictive maintenance and resource allocation across enterprise systems. The integration of ML techniques with process mining has fundamentally expanded capabilities beyond traditional descriptive analytics toward predictive and prescriptive applications that deliver substantial operational value. DL architectures including LSTM networks, transformers, and autoencoders have demonstrated exceptional performance in capturing complex temporal patterns and process behaviors from event log data. RL frameworks provide powerful approaches for learning optimal resource allocation policies through experience and environmental interaction.

Predictive maintenance applications leveraging ML-enhanced process mining enable organizations to transition from reactive approaches to proactive strategies that anticipate failures before disruptions occur. The combination of process event logs with sensor data creates rich multimodal datasets that support accurate predictions. Anomaly detection techniques complement direct failure prediction by identifying unusual patterns indicating emerging problems. Resource allocation optimization through DRL has shown remarkable success in discovering policies that substantially outperform traditional heuristics.

Several critical challenges must be addressed to realize the full potential of these technologies. Data quality issues significantly impact model performance and require substantial preprocessing effort. Model interpretability remains essential for building trust and enabling effective human oversight. Scalability to large-scale enterprise environments demands continued innovation in efficient algorithms and deployment architectures. Generalization across different organizational contexts requires advances in transfer learning and robust model development.

Future research directions include developing hybrid architectures that combine complementary strengths of different ML paradigms, advancing causal inference methods for process data, exploring federated learning for privacy-preserving model development, and investigating self-supervised learning techniques. The growing maturity of ML technologies and increasing availability of enterprise process data position ML-enhanced process mining as an increasingly critical capability for modern enterprises seeking to optimize operations and maintain competitive advantages through intelligent process management.

References

Vom Brocke, J., Jans, M., Mendling, J., & Reijers, H. A. (2021). A five-level framework for research on process mining. *Business & Information Systems Engineering*, *63*(5), 483–490. https://doi.org/10.1007/s12599-021-00698-8

Fahland, D., Ghidini, C., Becker, J., & Dumas, M. (2020). *Business process management*. Springer. https://doi.org/10.1007/978-3-030-49418-6

Carmona, J., van Dongen, B., & Weidlich, M. (2022). Conformance checking: Foundations, milestones and challenges. In W. M. P. van der Aalst & J. Carmona (Eds.), *Process mining handbook* (pp. 155–190). Springer. https://doi.org/10.1007/978-3-031-08848-0_7

Teinemaa, I., Dumas, M., La Rosa, M., & Maggi, F. M. (2019). Outcome-oriented predictive process monitoring: Review and benchmark. *ACM Transactions on Knowledge Discovery from Data*, *13*(2), Article 17. https://doi.org/10.1145/3301300

Fernandes, M., Corchado, J. M., & Marreiros, G. (2022). Machine learning techniques applied to mechanical fault diagnosis and fault prognosis in real industrial manufacturing use cases: A systematic literature review. *Applied Intelligence*, *52*(12), 14246–14280. https://doi.org/10.1007/s10489-022-03453-4

Altenmüller, T., Stüker, T., Waschneck, B., Kuhnle, A., & Lanza, G. (2020). Reinforcement learning for intelligent and autonomous production control of complex job-shops under time constraints. *Production Engineering, 14*(3), 319–328. https://doi.org/10.1007/s11740-020-00949-7

Rama-Maneiro, E., Vidal, J. C., & Lama, M. (2023). Deep learning for predictive business process monitoring: Review and benchmark. *IEEE Transactions on Services Computing, 16*(1), 739–756. https://doi.org/10.1109/TSC.2021.3063828

Neu, D. A., Lahann, J., & Fettke, P. (2022). A systematic literature review on state-of-the-art deep learning methods for process prediction. *Artificial Intelligence Review, 55*(2), 801–827. https://doi.org/10.1007/s10462-021-10030-3

Sivertsvik, M., Sumskiy, K., & Misimi, E. (2024, May). Learning active manipulation to target shapes with model-free, long-horizon deep reinforcement learning. In *Proceedings of the 2024 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5411–5418). IEEE. https://doi.org/10.1109/ICRA57147.2024.10611387

Kratsch, W., Manderscheid, J., Röglinger, M., & Seyfried, J. (2021). Machine learning in business process monitoring: A comparison of deep learning and classical approaches used for outcome prediction. *Business & Information Systems Engineering, 63*(3), 261–276. https://doi.org/10.1007/s12599-020-00662-4

Ter Hofstede, A. H. M., Koschmider, A., Marrella, A., Andrews, R., Fischer, D. A., Sadeghianasl, S., … Soffer, P. (2023). Process-data quality: The true frontier of process mining. *ACM Journal of Data and Information Quality, 15*(3), Article 18. https://doi.org/10.1145/3580254

Abou El Houda, Z., Brik, B., & Khoukhi, L. (2022). Why should I trust your IDS? An explainable deep learning framework for intrusion detection systems in IoT networks. *IEEE Open Journal of the Communications Society, 3*, 1164–1176. https://doi.org/10.1109/OJCOMS.2022.3188194

Lwakatare, L. E., Raj, A., Crnkovic, I., Bosch, J., & Olsson, H. H. (2020). Large-scale machine learning systems in real-world industrial settings: A review of challenges and solutions. *Information and Software Technology, 127*, 106368. https://doi.org/10.1016/j.infsof.2020.106368

Di Francescomarino, C. (2019). Predictive business process monitoring. In S. Sakr & A. Zomaya (Eds.), *Encyclopedia of Big Data Technologies* (pp. 1271–1280). Springer. https://doi.org/10.1007/978-3-319-63962-8_321-1

Mehdiyev, N., Fettke, P., & Evermann, J. (2020). A novel business process prediction model using a deep learning method. *Business & Information Systems Engineering, 62*(2), 143–157. https://doi.org/10.1007/s12599-019-00601-9

Verenich, I., Dumas, M., La Rosa, M., Maggi, F. M., & Teinemaa, I. (2019). Survey and cross-benchmark comparison of remaining time prediction methods in business process monitoring. *ACM Transactions on Intelligent Systems and Technology, 10*(4), Article 34. https://doi.org/10.1145/3312739

Jalayer, A., Kahani, M., Beheshti, A., Pourmasoumi, A., & Motahari-Nezhad, H. R. (2020, October). Attention mechanism in predictive business process monitoring. In *Proceedings of the IEEE 24th International Enterprise Distributed Object Computing Conference (EDOC)* (pp. 181–186). IEEE. https://doi.org/10.1109/EDOC49727.2020.00032

Nguyen, A., Chatterjee, S., Weinzierl, S., Schwinn, L., Matzner, M., & Eskofier, B. (2020). Time matters: Time-aware LSTMs for predictive business process monitoring. In *Proceedings of the International Conference on Process Mining* (pp. 112–123). Springer. https://doi.org/10.1007/978-3-030-37453-2_9

Camargo, M., Dumas, M., & González-Rojas, O. (2019). Learning accurate LSTM models of business processes. In *Proceedings of the International Conference on Business Process Management* (pp. 286–302). Springer. https://doi.org/10.1007/978-3-030-30493-5_17

Taymouri, F., La Rosa, M., Erfani, S. M., Bozorgi, Z. D., & Verenich, I. (2020). Predictive business process monitoring via generative adversarial nets: The case of next event prediction. In *Proceedings of the International Conference on Business Process Management* (pp. 237–256). Springer. https://doi.org/10.1007/978-3-030-58779-6_14

Carratù, M., Gallo, V., Iacono, S. D., Sommella, P., Bartolini, A., Grasso, F., … Patrizi, G. (2023). A novel methodology for unsupervised anomaly detection in industrial electrical systems. *IEEE Transactions on Instrumentation and Measurement, 72*, 1–12. https://doi.org/10.1109/TIM.2023.3241046

Zhang, W., Yang, D., & Wang, H. (2019). Data-driven methods for predictive maintenance of industrial equipment: A survey. *IEEE Systems Journal, 13*(3), 2213–2227. https://doi.org/10.1109/JSYST.2019.2896354

Ji, D., Wang, C., Li, J., & Dong, H. (2021). Data-driven-based fault diagnosis and remaining useful life prediction of petroleum machinery and equipment: A review. *Systems Science & Control Engineering, 9*(1), 724–747. https://doi.org/10.1080/21642583.2021.1964958

De Gasperis, G., & Facchini, S. D. (2025). A comparative study of rule-based and data-driven approaches in industrial monitoring. *arXiv preprint* arXiv:2509.15848.

Huo, S., Völzer, H., Reddy, P., Agarwal, P., Isahagian, V., & Muthusamy, V. (2021). Graph autoencoders for business process anomaly detection. In *Proceedings of the International Conference on Business Process Management* (pp. 417–433). Springer. https://doi.org/10.1007/978-3-030-85469-0_25

Ko, J., & Comuzzi, M. (2023). A systematic review of anomaly detection for business process event logs. *Business & Information Systems Engineering, 65*(4), 441–462. https://doi.org/10.1007/s12599-022-00755-9

Pachauri, P. (2025). A hybrid deep learning framework for anomaly detection in high-dimensional streaming data: Integrating autoencoders and LSTM networks. *JANOLI International Journal of Big Data, 1*(2).

Neubauer, T. R., da Silva, V. F., Fantinato, M., & Peres, S. M. (2022, November). Resource allocation optimization in business processes supported by reinforcement learning and process mining. In *Proceedings of the Brazilian Conference on Intelligent Systems* (pp. 580–595). Springer. https://doi.org/10.1007/978-3-031-21689-0_43

Arias, M., Rojas, E., Aguirre, S., Cornejo, F., Muñoz-Gama, J., Sepúlveda, M., & Capurro, D. (2020). Mapping the patient's journey in healthcare through process mining. *International Journal of Environmental Research and Public Health, 17*(18), 6586. https://doi.org/10.3390/ijerph17186586

Fayaz, S. A., Jahangeer Sidiq, S., Zaman, M., & Butt, M. A. (2022). Machine learning: An introduction to reinforcement learning. In *Machine learning and data science: Fundamentals and applications* (pp. 1–22). Springer.

Tran-Dang, H., Bhardwaj, S., Rahim, T., Musaddiq, A., & Kim, D. S. (2022). Reinforcement learning-based resource management for fog computing environments: A literature review, challenges, and open issues. *Journal of Communications and Networks, 24*(1), 83–98. https://doi.org/10.1109/JCN.2022.000004

Pang, G., Wang, X., Wang, L., Hao, F., Lin, Y., Wan, P., & Min, G. (2022). Efficient deep reinforcement learning-enabled recommendation. *IEEE Transactions on Network Science and Engineering, 10*(2), 871–886. https://doi.org/10.1109/TNSE.2022.3164219

Rizzi, W., Di Francescomarino, C., & Maggi, F. M. (2020). Explainability in predictive process monitoring: When understanding helps improving. In *Proceedings of the International Conference on Business Process Management* (pp. 141–158). Springer. https://doi.org/10.1007/978-3-030-58666-9_9

Zhao, L., Li, H., Wang, J., & Chen, L. (2025). Timing and resource analysis of emergency response processes using Petri nets. *IEEE Access, 13*, 1–15.

Yeon, M. S., Lee, Y. K., Pham, D. L., & Kim, K. P. (2022). Experimental verification of human-centric network-based resource allocation approaches for process-aware information systems. *IEEE Access, 10*, 23342–23354. https://doi.org/10.1109/ACCESS.2022.3158169

Park, G., & Song, M. (2020). Predicting performances in business processes using deep neural networks. *Decision Support Systems, 129*, 113191. https://doi.org/10.1016/j.dss.2019.113191

Nashaat, M., & Miller, J. (2024). Towards efficient fine-tuning of language models with organizational data for automated software review. *IEEE Transactions on Software Engineering*. Advance online publication. https://doi.org/10.1109/TSE.2024.3358123

Galanti, R., Coma-Puig, B., de Leoni, M., Carmona, J., & Navarin, N. (2020). Explainable predictive process monitoring. In *Proceedings of the 2nd International Conference on Process Mining* (pp. 1–8). IEEE. https://doi.org/10.1109/ICPM49681.2020.00007

Zintgraf, L., Shiarli, K., Kurin, V., Hofmann, K., & Whiteson, S. (2019). Fast context adaptation via meta-learning. In *Proceedings of the International Conference on Machine Learning* (pp. 7693–7702). PMLR.

Nannini, L., Bolos, A. C., Lama, M., & Barro, S. (2023). Explainable AI in process mining: A systematic literature review. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5162771

Wang, J., Wiens, J., & Lundberg, S. (2021). Shapley flow: A graph-based approach to interpreting model predictions. In *Proceedings of the International Conference on Artificial Intelligence and Statistics* (pp. 721–729). PMLR.

Turek, J., Jain, S., Vo, V., Capotă, M., Huth, A., & Willke, T. (2020). Approximating stacked and bidirectional recurrent architectures with the delayed recurrent neural network. In *Proceedings of the International Conference on Machine Learning* (pp. 9648–9658). PMLR.

Chandar, S., Sankar, C., Vorontsov, E., Kahou, S. E., & Bengio, Y. (2019). Towards non-saturating recurrent units for modelling long-term dependencies. In *Proceedings of the AAAI Conference on Artificial Intelligence, 33*(1), 3280–3287. https://doi.org/10.1609/aaai.v33i01.33013280

Beck, M., Pöppel, K., Spanring, M., Auer, A., Prudnikova, O., Kopp, M., … Hochreiter, S. (2024). xLSTM: Extended long short-term memory. *Advances in Neural Information Processing Systems, 37*, Article 107547.

Yu, Y., Si, X., Hu, C., & Zhang, J. (2019). A review of recurrent neural networks: LSTM cells and network architectures. *Neural Computation, 31*(7), 1235–1270. https://doi.org/10.1162/neco_a_01199

Shewalkar, A. (2019). Performance evaluation of deep neural networks applied to speech recognition: RNN, LSTM and GRU. *Journal of Artificial Intelligence and Soft Computing Research, 9*(4), 235–246. https://doi.org/10.2478/jaiscr-2019-0018

Zhang, J., Karimireddy, S. P., Veit, A., Kim, S., Reddi, S. J., Kumar, S., & Sra, S. (2020). Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems, 33*, 15383–15393.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT* (pp. 4171–4186). Association for Computational Linguistics. https://doi.org/10.18653/v1/N19-1423

Camargo, M., Dumas, M., & González-Rojas, O. (2021). Discovering generative models from event logs: Data-driven simulation vs deep learning. *PeerJ Computer Science, 7*, e577. https://doi.org/10.7717/peerj-cs.577

Montesinos López, O. A., Montesinos López, A., & Crossa, J. (2022). Fundamentals of artificial neural networks and deep learning. In *Multivariate statistical machine learning methods for genomic prediction* (pp. 379–425). Springer. https://doi.org/10.1007/978-3-030-89010-0_14

asquadibisceglie, V., Appice, A., Castellano, G., & Malerba, D. (2019, June). Using convolutional neural networks for predictive process analytics. In *Proceedings of the 2019 IEEE International Conference on Process Mining (ICPM)* (pp. 129–136). IEEE. https://doi.org/10.1109/ICPM.2019.00029

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.

Khemakhem, I., Kingma, D. P., Monti, R. P., & Hyvärinen, A. (2020). Variational autoencoders and nonlinear ICA: A unifying framework. In *Proceedings of the International Conference on Artificial Intelligence and Statistics (AISTATS)* (pp. 2207–2217). PMLR.

Scarselli, F., Gori, M., Tsoi, A. C., Hagenbuchner, M., & Monfardini, G. (2009). The graph neural network model. *IEEE Transactions on Neural Networks, 20*(1), 61–80. https://doi.org/10.1109/TNN.2008.2005605

Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., & Yu, P. S. (2020). A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems, 32*(1), 4–24. https://doi.org/10.1109/TNNLS.2020.2978386

Dalzochio, J., Kunst, R., Pignaton, E., Binotto, A., Sanyal, S., Favilla, J., & Barbosa, J. (2020). Machine learning and reasoning for predictive maintenance in Industry 4.0: Current status and challenges. *Computers in Industry, 123*, 103298. https://doi.org/10.1016/j.compind.2020.103298

Morettini, S. (2021). *Machine learning in predictive maintenance of industrial robots* (Master's thesis). University of Bologna.

Zhang, L., Lin, J., Liu, B., Zhang, Z., Yan, X., & Wei, M. (2019). A review on deep learning applications in prognostics and health management. *IEEE Access, 7*, 162415–162438. https://doi.org/10.1109/ACCESS.2019.2950362

Wu, F., Wu, Q., Tan, Y., & Xu, X. (2024). Remaining useful life prediction based on deep learning: A survey. *Sensors, 24*(11), 3454. https://doi.org/10.3390/s24113454

Zhou, J., Yang, J., Qian, Q., & Qin, Y. (2024). A comprehensive survey of machine remaining useful life prediction approaches based on pattern recognition: Taxonomy and challenges. *Measurement Science and Technology, 35*(6), 062001. https://doi.org/10.1088/1361-6501/ad2b3a

Ergen, T., & Kozat, S. S. (2019). Unsupervised anomaly detection with LSTM neural networks. *IEEE Transactions on Neural Networks and Learning Systems, 31*(8), 3127–3141. https://doi.org/10.1109/TNNLS.2019.2943735

Audibert, J., Michiardi, P., Guyard, F., Marti, S., & Zuluaga, M. A. (2020). USAD: Unsupervised anomaly detection on multivariate time series. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* (pp. 3395–3404). ACM. https://doi.org/10.1145/3394486.3403392

Feng, Y., Cai, W., Yue, H., Xu, J., Lin, Y., Chen, J., & Hu, Z. (2022). An improved X-means and isolation forest–based methodology for network traffic anomaly detection. *PLOS ONE, 17*(1), e0263423. https://doi.org/10.1371/journal.pone.0263423

Pasquadibisceglie, V., Appice, A., Castellano, G., & Malerba, D. (2022). A multi-view deep learning approach for predictive business process monitoring. *IEEE Transactions on Services Computing, 15*(4), 2382–2395. https://doi.org/10.1109/TSC.2021.3050106

Schonlau, M., & Zou, R. Y. (2020). The random forest algorithm for statistical learning. *The Stata Journal, 20*(1), 3–29. https://doi.org/10.1177/1536867X20909688

Cherif, I. L., & Kortebi, A. (2019). On using extreme gradient boosting (XGBoost) machine learning algorithm for home network traffic classification. In *Proceedings of 2019 Wireless Days (WD)* (pp. 1–6). IEEE. https://doi.org/10.1109/WD.2019.8734194

Chen, Z., Chen, L., Shen, W., & Xu, K. (2022). Remaining useful life prediction of lithium-ion batteries via sequence decomposition and deep learning integrated approach. *IEEE Transactions on Vehicular Technology, 71*(2), 1466–1479. https://doi.org/10.1109/TVT.2021.3134814

Tarawneh, A. S., Hassanat, A. B., Almohammadi, K., Chetverikov, D., & Bellinger, C. (2020). SMOTEFUNA: Synthetic minority over-sampling technique based on furthest neighbour algorithm. *IEEE Access, 8*, 59069–59082. https://doi.org/10.1109/ACCESS.2020.2982927

Sterner, P., Goretzko, D., & Pargent, F. (2025). Everything has its price: Foundations of cost-sensitive machine learning and its application in psychology. *Psychological Methods, 30*(1), 112–129. https://doi.org/10.1037/met0000563

Liu, Z., Li, Y., Chen, N., Wang, Q., Hooi, B., & He, B. (2025). A survey of imbalanced learning on graphs: Problems, techniques, and future directions. *IEEE Transactions on Knowledge and Data Engineering*. Advance online publication.

Pech, M., Vrchota, J., & Bednář, J. (2021). Predictive maintenance and intelligent sensors in smart factory. *Sensors, 21*(4), 1470. https://doi.org/10.3390/s21041470

He, J., Mao, R., Shao, Z., & Zhu, F. (2020). Incremental learning in online scenario. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 13926–13935). IEEE.

Qiu, T., Chi, J., Zhou, X., Ning, Z., Atiquzzaman, M., & Wu, D. O. (2020). Edge computing in industrial internet of things: Architecture, advances, and challenges. *IEEE Communications Surveys & Tutorials, 22*(4), 2462–2488. https://doi.org/10.1109/COMST.2020.2981391

Harl, M., Weinzierl, S., Stierle, M., & Matzner, M. (2020). Explainable predictive business process monitoring using gated graph neural networks. *Journal of Decision Systems, 29*(sup1), 312–327. https://doi.org/10.1080/12460125.2020.1768524

Zhang, C., Song, D., Chen, Y., Feng, X., Lumezanu, C., Cheng, W., & Chawla, N. V. (2019). A deep neural network for unsupervised anomaly detection and diagnosis in multivariate time series data. In *Proceedings of the AAAI Conference on Artificial Intelligence, 33*(1), 1409–1416. https://doi.org/10.1609/aaai.v33i01.33011409

Kuru, K. (2024). *Towards state and situation awareness for driverless vehicles using deep neural networks* (Technical report).

Zahavy, T., Xu, Z., Veeriah, V., Hessel, M., Oh, J., van Hasselt, H., … Singh, S. (2020). A self-tuning actor-critic algorithm. *Advances in Neural Information Processing Systems, 33*, 20913–20924.

Li, X., & Cohen, K. (2020). Improved fuzzy Q-learning with replay memory. In *Proceedings of the North American Fuzzy Information Processing Society Annual Conference* (pp. 13–23). Springer. https://doi.org/10.1007/978-3-030-38570-5_2

Janisch, J., Pevný, T., & Lisý, V. (2019). Classification with costly features using deep reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence, 33*(1), 3959–3966. https://doi.org/10.1609/aaai.v33i01.33013959

Wang, Y., He, H., & Tan, X. (2020). Truly proximal policy optimization. In *Proceedings of the Uncertainty in Artificial Intelligence (UAI) Conference* (pp. 113–122). PMLR.

Böhm, P., Pounds, P., & Chapman, A. C. (2022). Non-blocking asynchronous training for reinforcement learning in real-world environments. In *Proceedings of the 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 10927–10934). IEEE. https://doi.org/10.1109/IROS56902.2022.9981195

Turner, C. J., Ma, R., Chen, J., & Oyekan, J. (2021). Human in the loop: Industry 4.0 technologies and scenarios for worker mediation of automated manufacturing. *IEEE Access, 9*, 103950–103966. https://doi.org/10.1109/ACCESS.2021.3098470

Nguyen, T. T., Nguyen, N. D., & Nahavandi, S. (2020). Deep reinforcement learning for multiagent systems: A review of challenges, solutions, and applications. *IEEE Transactions on Cybernetics, 50*(9), 3826–3839. https://doi.org/10.1109/TCYB.2020.2977613

Shao, K., Zhu, Y., Tang, Z., & Zhao, D. (2020). Cooperative multi-agent deep reinforcement learning with counterfactual reward. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)* (pp. 1–8). IEEE. https://doi.org/10.1109/IJCNN48605.2020.9207090

Zhou, M., Liu, Z., Sui, P., Li, Y., & Chung, Y. Y. (2020). Learning implicit credit assignment for cooperative multi-agent reinforcement learning. *Advances in Neural Information Processing Systems, 33*, 11853–11864.

Hernandez-Leal, P., Kartal, B., & Taylor, M. E. (2019). A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems, 33*(6), 750–797. https://doi.org/10.1007/s10458-019-09421-1

Boukredera, D., & Adel-Aissanou, K. (2020). Modeling and performance analysis of cognitive radio networks using stochastic timed colored Petri nets. *Wireless Personal Communications, 112*(3), 1659–1687. https://doi.org/10.1007/s11277-020-07147-5

Shah, S. H., & Duni, F. (2021). A review on matrix factorization techniques used for an intelligent recommender system. *Turkish Journal of Computer and Mathematics Education, 12*(7), 1812–1823.

Wissenburg, R., Kusters, R. J., & Martin, H. (2023). IT project portfolio management: Development and validation of a reference model. In *Proceedings of the International Conference on Enterprise Information Systems (ICEIS)* (Vol. 1, pp. 202–211). SCITEPRESS.

Wagner, G. (2021). Business process modeling and simulation with DPMN: Processing activities. In *Proceedings of the 2021 Winter Simulation Conference (WSC)* (pp. 1–15). IEEE. https://doi.org/10.1109/WSC52266.2021.9715458

Dai, J., Ji, J., Yang, L., Zheng, Q., & Pan, G. (2023). Augmented proximal policy optimization for safe reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence, 37*(6), 7288–7295. https://doi.org/10.1609/aaai.v37i6.25971

Altman, E. (2021). *Constrained Markov decision processes.* Routledge.

Carvalho, D. S., Melo, F. S., & Santos, P. A. (2023). Theoretical remarks on feudal hierarchies and reinforcement learning. In *Proceedings of ECAI 2023* (pp. 351–356). IOS Press.

Liu, Y., Ding, J., & Liu, X. (2020). IPO: Interior-point policy optimization under constraints. In *Proceedings of the AAAI Conference on Artificial Intelligence, 34*(4), 4940–4947. https://doi.org/10.1609/aaai.v34i04.5881

van Cruchten, R. (2020). Data quality in process mining: A rule-based approach. In *Interactive Process Mining in Healthcare* (pp. 53–79). Springer.

Acitelli, G., Angelini, M., Bonomi, S., Maggi, F. M., Marrella, A., & Palma, A. (2022). Context-aware trace alignment with automated planning. In *Proceedings of the 4th International Conference on Process Mining (ICPM)* (pp. 104–111). IEEE. https://doi.org/10.1109/ICPM57379.2022.00020

Wang, M., Zhang, X., Yang, Y., & Wang, J. (2025). Explainable machine learning in risk management: Balancing accuracy and interpretability. *Journal of Financial Risk Management, 14*(3), 185–198.

Zhang, S., Qiu, L., & Zhang, H. (2025). Edge cloud synergy models for ultra-low latency data processing in smart city IoT networks. *International Journal of Science, 12*(10).

Yang, J., Zeng, Z., & Shen, Z. (2025). Neural-symbolic dual-indexing architectures for scalable retrieval-augmented generation. *IEEE Access.* Advance online publication.

Molnar, C. (2020). *Interpretable machine learning* (2nd ed.). Lulu.com.

Ben-Nun, T., & Hoefler, T. (2019). Demystifying parallel and distributed deep learning: An in-depth concurrency analysis. *ACM Computing Surveys, 52*(4), Article 65. https://doi.org/10.1145/3320060

Choudhary, T., Mishra, V., Goswami, A., & Sarangapani, J. (2020). A comprehensive survey on model compression and acceleration. *Artificial Intelligence Review, 53*(7), 5113–5155. https://doi.org/10.1007/s10462-020-09816-7

Causal-aware multimodal transformer for supply chain demand forecasting: Integrating text, time series, and satellite imagery. (n.d.). Unpublished manuscript.

Becker, A., & Becker, J. (2021). Dataset shift assessment measures in monitoring predictive models. *Procedia Computer Science, 192*, 3391–3402. https://doi.org/10.1016/j.procs.2021.09.112

Hosna, A., Merry, E., Gyalmo, J., Alom, Z., Aung, Z., & Azim, M. A. (2022). Transfer learning: A friendly introduction. *Journal of Big Data, 9*(1), 102. https://doi.org/10.1186/s40537-022-00689-7

Niu, S., Liu, Y., Wang, J., & Song, H. (2020). A decade survey of transfer learning (2010–2020). *IEEE Transactions on Artificial Intelligence, 1*(2), 151–166. https://doi.org/10.1109/TAI.2020.3019323

Vom Brocke, J., Baier, M. S., Schmiedel, T., Stelzl, K., Röglinger, M., & Wehking, C. (2021). Context-aware business process management: Method assessment and selection. *Business & Information Systems Engineering, 63*(5), 533–550. https://doi.org/10.1007/s12599-021-00690-2

Bansal, G., Nushi, B., Kamar, E., Lasecki, W. S., Weld, D. S., & Horvitz, E. (2019). Beyond accuracy: The role of mental models in human–AI team performance. In *Proceedings of the AAAI Conference on Human Computation and Crowdsourcing, 7*, 2–11. https://doi.org/10.1609/hcomp.v7i1.5275

Ren, S., Chen, S., & Zhang, Q. (2025). Reinforcement learning paradigms for proactive cybersecurity and dynamic risk management. *Frontiers in Artificial Intelligence Research, 2*(3), 436–456.

Mai, N. T., Cao, W., & Fang, Q. (2025). A study on how LLMs (e.g., GPT-4, chatbots) are being integrated to support tutoring, essay feedback, and content generation. *Journal of Computing and Electronic Information Management, 18*(3), 43–52.