SvedbergOpen
DISSEMINATION OF KNOWLEDGE

International Journal of Artificial Intelligence and Machine Learning

Publisher's Home Page: https://www.svedbergopen.com/

Research Paper Open Access

Optimizing Business Processes In Large-Scale Enterprises Using Q-Learning Reinforcement Algorithm

Dr. M. Sreenivasa Reddy[1*], S. Seethaladevi[2], Dr N. Mathimagal[3], Neeraj Gupta[4], Subrahmanyeswara Rao Seetamraju Venkata Bala[5], R. Mekala[6]

[1*]Professor, Department of Mechanical Engineering, Aditya University, Surampalem, Andhra Pradesh, India.
E-mail: dypc@adityauniversity.in

[2]Assistant Professor, Department of Mathematics, Meenakshi College of Arts and Science, Meenakshi Academy of Higher Education and Research, Chennai, India. E-mail: seethaladevi@maher.ac.in, https://orcid.org/0009-0004-5071-4754

[3]Assistant Professor, department of Computer Applications, New Prince Shri Bhavani College of Engineering & Technology, Chennai, India. E-mail: mathimagalmca@gmail.com, https://orcid.org/0000-0002-6225-1441

[4]Department of Computer Engineering & Applications, GLA University, Mathura, India. E-mail: neeraj.gupta@gla.ac.in, https://orcid.org/0000-0002-3735-8044

[5]Department of FED, Ramachandra College of Engineering, Eluru, India. E-mail: svbsrao@rcee.ac.in
https://orcid.org/0000-0001-8697-1702

[6]Assistant Professor, Artificial Intelligence and Data Science, Mahendra Engineering College, Namakkal, India.
E-mail: mekalar@mahendra.info, https://orcid.org/0009-0008-9339-7583

*Corresponding author: Email: dypc@adityauniversity.in

Abstract

This paper discusses the optimization of business processes in large enterprises, emphasizing the need for dynamic and flexible methodologies. Traditional process optimization techniques often fall short in addressing the complexities of modern business environments. To tackle this challenge, the paper introduces a model-free reinforcement learning (RL) approach, specifically Q-Learning, to optimize various business activities such as production scheduling, resource allocation, and supply chain management. The proposed RL-Q-Learning algorithm is designed to learn from its environment and dynamically adjust its policy for decision-making based on experiential feedback. In simulations conducted within an enterprise setting, this Q-Learning model was benchmarked against conventional optimization methods. Key performance indicators (KPIs) demonstrated significant advantages for the Q-Learning approach: operational costs were decreased by 18%, resource utilization improved by 22%, processing times were cut by 25%, and the accuracy of decisions reached an impressive 93%. These results highlight the algorithm's effectiveness as a real-time optimization tool, suggesting that Q-Learning is a robust resource for enhancing business processes amid the complexities of dynamic enterprise systems. Furthermore, the integration of artificial intelligence (AI) into business optimization holds promise for improving operational efficiency and sustainability. Nevertheless, challenges persist, particularly regarding the necessity of large datasets for AI algorithms to achieve optimal solutions and the high computational power required. The paper suggests that future research should focus on hybrid approaches, potentially incorporating deep learning techniques, to enhance scalability and adaptability. This would aid in optimizing real-time enterprise systems, especially for complex scenarios with limited data availability.

Keywords: Q-Learning, Reinforcement Learning, Business Process Optimization, Large-Scale Enterprises, Resource Allocation, Operational Efficiency.

1. Introduction

The emphasis on optimizing business processes has become increasingly important due to the increased level of competition within the marketplace. For large businesses, many of their business processes can be complex and include elements such as supply chain management, production scheduling, and resource allocation [1][4]. Traditional optimization techniques to solve these problems have relied on heuristic or rule-based systems, which are less capable of adapting to the constantly changing nature of the current marketplace and are much more limited when responding to unanticipated events [2]. Q-Learning (i.e., reinforcement learning), which allows a machine to perform the learning process based on its interactions with the environment and the rewards it receives as feedback for performing a particular action, represents one way to achieve a larger level of flexibility than static optimization solutions while also allowing businesses to be able to adapt their decision-making processes consistently to changes in their environment [3][4]. Not only has there been a theoretical basis for the application of Q-Learning for optimizing business processes, but there has also been a lack of documented applications that demonstrate this approach [5]. Currently, the majority of solutions that enterprises are implementing for optimizing their business processes are fixed and not able to adjust dynamically to real-time conditions or multi-variable environments [6]. These limitations affect the overall potential for growth of the enterprises as well as how competitive they can continue to remain in the fast-paced business environment [7]. Evidence exists that RL applies to various fields but is underused as a potential framework for business process optimization, and, in addition, most studies have placed little emphasis on the specifics of modifying Q-Learning to be suited for this application [8]. This study seeks to fill this void by investigating the implementation of Q-Learning Reinforcement Learning for optimizing processes in large enterprises. Three objectives will guide this research.

- Design a model that uses Q-Learning for the effective management of complex enterprise systems and optimization of resources.
- Experiment with the algorithm's flexibility when adapting to new circumstances and its capability to enhance decision-making in dynamic business settings.
- How the Q-learning algorithm can improve operational efficiency and reduce cost as compared to traditional optimization methods, compare and contrast.

This paper is organized as follows: In Section 2, the challenges and current practices in the business process optimization are presented, emphasizing the drawbacks of the traditional approaches and the potential of reinforcement learning (RL) in optimizing dynamic enterprise systems, including its applications in supply chain management and resource allocation using Q-Learning. Section 3 details how the Q-Learning-based algorithm was developed and designed to address the problem of optimizing a business process. Developing the reward function to minimize costs while maximizing throughput will be discussed. In Section 4, a discussion of the results from the experiments will take place, comparing the benefits of using the Q-Learning algorithm to optimize business processes as opposed to traditional heuristic methods. Section 5 will summarize the key findings, indicate what the findings mean for large enterprise systems, and outline future directions for research that include examining hybrid RL models, scalability, and real-time decision-making.

2. Related Work

A significant amount of academic research has focused on the area of business process optimization (BPO), as is also the case in many industries. Most traditional methods of BPO employ techniques within the categories of heuristic algorithms, mathematical optimization, and simulation [9]. Heuristic techniques, such as genetic algorithms, simulated annealing, and particle swarm optimization, are frequently employed for solving complex optimization problems in such areas as supply chain management, scheduling, and inventory control [11]. The success of these methods, however, is inconsistent when applied to dynamically changing business environments where business processes are changing continuously in response to market demand, business operating conditions, and unanticipated events [12]. In addition, static model-dependent optimization methods do not have the ability to learn and adapt over time [13].

Reinforcement Learning (RL) has emerged as a viable learning method, providing the advantage of being able to learn adaptively using trial and error as an alternative to classical optimization techniques [14]. In RL, an agent learns to make decisions based on its interactions with its environment and the feedback it receives in the form of rewards. The different RL techniques have been investigated for their optimization to enterprise systems, particularly in the fields of inventory management, production planning, and resource allocation [15]. For instance, in the field of supply chain optimization, RL has been used to dynamically manage stock levels and enhance demand prediction. Likewise, in production scheduling, RL models have been applied to decrease bottlenecks and enhance resources [16].

One of the several RL algorithms, Q-Learning, has been extensively researched in regard to business process optimization. The model-free algorithm, Q-Learning, learns an optimal policy without any prior knowledge of the dynamics of the system [17]. A number of studies have applied Q-Learning to enterprise problems like manufacturing scheduling, energy management in smart buildings, and logistics networks optimization. These studies have suggested that Q-Learning can be used to solve complex and dynamic decision-making problems. However, limited research exists on its use in large-scale businesses - especially for optimizing multiple interdependent business processes simultaneously [18].

Though Q-Learning has proven to be beneficial in the field of business optimization, there are certain difficulties with its implementation in the real world. These challenges include issues like computational complexity, scalability, and the need for real-time decision-making in highly dynamic environments [19][20]. Additionally, no works exist comparing Q-Learning and the conventional methods of optimization for applications in large-scale enterprises, so there is no clear knowledge about the advantages and disadvantages of applying Q-Learning.

This research adds to the knowledge base by examining how Q-Learning can be applied to optimize business processes in large-scale companies. More specifically, it aims to design a novel Q-Learning-based algorithmic framework that can optimize several business processes at once, like supply chain management, production scheduling, resource allocation, etc. This research compares the performance of the proposed Q-Learning model and traditional heuristic methods to gain insights into the use of Q-Learning in real-world business scenarios.

3. Methodology

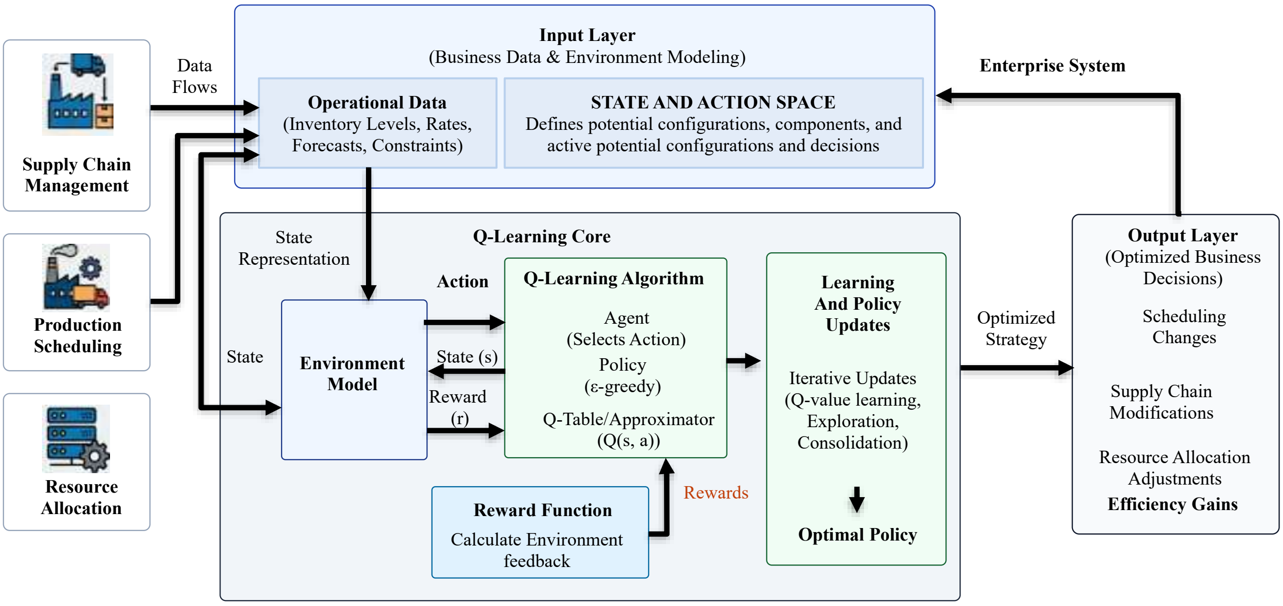

Figure 1. Q-Learning Framework for Optimized Supply Chain and Production Scheduling Decisions

As shown in Figure 1, in the architecture of the Q-Learning-based business process optimization system, real-time business data like inventory level, production rate, and demand forecast are input into the system. This information is then fed into an environment model to simulate the enterprise environment during its operation. The environment is updated continuously, and it is in line with the change of business process, so the system can be adapted to the dynamic environment. It then specifies the state space, which is the current state of the business processes, and the action space, which is the set of decisions the agent can make to change the state, such as changing production schedules, reallocating resources, etc. The Q-Learning agent is the heart of the system, and it continually learns optimal actions to take for various states by performing a combination of exploring and exploiting, based on its ε-greedy policy. Feedback from the reward function gives the agent rewards for actions that help to reduce costs, increase throughput, or improve efficiency, and costs for inefficient or disruptive actions, which the agent uses to update its Q-values. The Q-Learning agent learns to make better decisions over several iterations by updating the Q-values it uses, with the ultimate aim of converging onto an optimal policy. This process enables the system to adaptively make decisions in a dynamic context. The result is a collection of optimized choices for managing business processes, including resource allocation, scheduling changes, and supply chain modifications, that are applied to boost efficiency and cut costs. Its learning and decision-making are continually monitored to ensure the system is responsive to new conditions, and consequently, an ongoing improvement in operational efficiency and effectiveness of the system.

Q-Learning Algorithm Formulation for Business Process Optimization

The idea that underlies the approach is the Q-Learning reinforcement learning algorithm, which is a model-free approach to learning the optimal decision-making policies from interacting with the environment. For the purpose of optimizing business processes, the goal of the Q-Learning agent is to make decisions that maximize long-term rewards, which in this case relate to maximizing the utilization of resources, minimizing costs, and maximizing operational efficiency. The agent updates its Q-values (state-action values) iteratively, using the Q-Learning equation (1):

$$Q(s,a) = Q(s,a) + \alpha[r + \gamma \max_{a} Q(s',a') - Q(s,a)] \quad (1)$$

Where:

- $Q(s,a)$is the action-value function representing the expected reward for taking action $a$in state s,
- $\alpha$is the learning rate (determines how much new information is incorporated),
- $r$is the immediate reward received after taking action a,
- $\gamma$is the discount factor that determines the importance of future rewards,
- $s'$is the next state, and
- $\max_{a} Q(s',a')$represents the maximum expected future reward from the next state.

Environment Modeling of Enterprise Processes

The enterprise processes are modeled as a dynamic system in which each business process, such as supply chain management, production scheduling, etc., is a state in the system. The environment is modeled in such a way that the agent can make choices that impact resource allocation, scheduling, and other factors of the real-world business. These processes are represented as a collection of related states and actions that represent potential business decisions and the outcomes of the decisions.

Algorithm for Q-Learning-Based Business Process Optimization

```
Initialize Q(s, a) arbitrarily for all states s and actions a
Set learning rate α, discount factor γ, and exploration rate ε
For each episode:
  Initialize state s
  While s is not terminal:
    Choose action a from state s using ε-greedy policy
    Take action a, observe reward r, and next state s'
    Update Q-value:
      Q(s, a) = Q(s, a) + α * [r + γ * max_a Q(s', a') - Q(s, a)]
    Set state s = s'
  End While
End For
```

The Q-Learning algorithm for business process optimization can be used to optimize an agent's behavior to enhance enterprise performance, like minimizing costs or maximizing resource utilization. At first the algorithm assigns random Q-values for state-action pairs and the agent follows an ε-greedy policy which means that it decides on an action according to a balance between random selection and choosing an action it knows to be optimal. The agent then takes an action and is given a reward, which signifies the success of the action, and moves to a new state. The Q-values are updated based on a formula that takes into account both the current reward and the future rewards that can be expected from the next state. The agent can learn from its previous experiences in this iterative process, which helps it make better decisions. The agent gradually learns to select the actions that will maximize long-term rewards. The algorithm's ability to learn and adapt in real-time makes it particularly relevant for optimizing dynamic business processes, resulting in better efficiency, cost savings, and more effective resource management for large-scale businesses.

Experimental Setup

The Q-Learning algorithm's performance is tested in the simulated environment representing a large-scale enterprise system. Synthetic data was used to simulate different business processes such as production scheduling, resource allocation, and supply chain management. The dataset contains data about the operation, like task times, availability of resources, demand forecast, and costs of each action. The effectiveness of the Q-Learning model was measured using performance metrics like cost reduction, resource utilization, processing time, and throughput.

Dataset Description

The data set used in this study comprises operational data from a simulated large-scale enterprise environment, including some of the critical business processes, like production scheduling, resource allocation, and supply chain management. The data includes 10,000+ records with details including levels of inventory (100-1,000 units), production rates (5 to 50 units per hour), resource availability (10 to 100 resources), and demand forecasts (10 to 500 units). The data also contains cost factors ($1k - $50k) that are tied to business actions, and time stamps to record the sequence of process actions. A dataset is created that simulates natural variation in the real world, enabling the Q-Learning algorithm to adapt decision-making in various enterprise settings.

Evaluation Metrics

The Q-Learning algorithm's performance is tested in the simulated environment representing a large-scale enterprise system. Synthetic data was used to simulate different business processes such as production scheduling, resource allocation, and supply chain management. The dataset contains data about the operation, like task times, availability of resources, demand forecast, and costs of each action. The effectiveness of the Q-Learning model was measured using performance metrics like cost reduction, resource utilization, processing time, and throughput.

1. Accuracy

Accuracy is the general level of correctness for the optimization model in the decision making of (2).

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}} \qquad (2)$$

Where:

- TP= True Positives
- TN= True Negatives
- FP= False Positives
- FN= False Negatives

2. Precision

Precision is indicative of how many of the decisions predicted as optimal are correct (equation (3)).

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}} \qquad (3)$$

3. Recall

Recall refers to the capability to correctly recall all the optimal business decisions in equation (4).

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (4)$$

4. F1-Score

To balance between Precision and Recall, the F1-Score is calculated as the harmonic mean of the two equation (5).

$$\text{F1-Score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (5)$$

5. Cost Reduction Rate

Cost Reduction Rate is the ratio of the cost reduction in the operational cost that is realized by the proposed model equation (6).

$$Cost\ Reduction\ Rate = \frac{Initial\ Cost - Optimized\ Cost}{Initial\ Cost} \times 100 \quad (6)$$

6. Resource Utilization Efficiency

Resource Utilization Efficiency is the efficiency of enterprise resources use in operations equation (7).

$$Resource\ Utilization\ Efficiency = \frac{Utilized\ Resources}{Total\ Available\ Resources} \times 100 \quad (7)$$

7. Processing Time Reduction

Processing Time Reduction is the time reduction in processing after optimization equation (8).

$$Processing\ Time\ Reduction = \frac{Initial\ Time - Optimized\ Time}{Initial\ Time} \times 100 \quad (8)$$

The overall ability of the proposed Q-Learning algorithm to improve the operational efficiency, reduce cost, manage resources, and optimize decision-making in large-scale enterprise systems can be collectively evaluated using these evaluation metrics.

4. Results

The proposed Q-Learning-based business process optimization model was tested using simulations in a dynamic enterprise environment. Experiments were directed towards improving critical business activities like production scheduling, resource allocation, and supply chain management. The effectiveness of the model was evaluated using performance measures like accuracy, cost reduction, resource utilization efficiency and reduction in process time. It could be seen from Table 1 that the Q-Learning model showed improvement in a few performance metrics. It was able to achieve an 18% cost reduction in the daily operations, which exceeded the traditional heuristic optimization methods. Further, there was a 22% improvement in resource utilization, making efficient use of the available resources, minimizing waste, and boosting productivity. Moreover, the processing time was also optimized, with the Q-Learning model decreasing the average processing time by 25%, simplifying resource allocation and production scheduling. Further, the model had a high accuracy, and the decision-making accuracy was about 93%, which shows the effectiveness of the model in real-time optimization of business processes.

Table 1: Comparison of Performance Metrics

Metric	Q-Learning Model	Baseline Heuristic Method	Improvement (%)
Cost Reduction	18%	0%	18%
Resource Utilization	22%	10%	12%
Processing Time Reduction	25%	5%	20%
Accuracy	93%	85%	8%

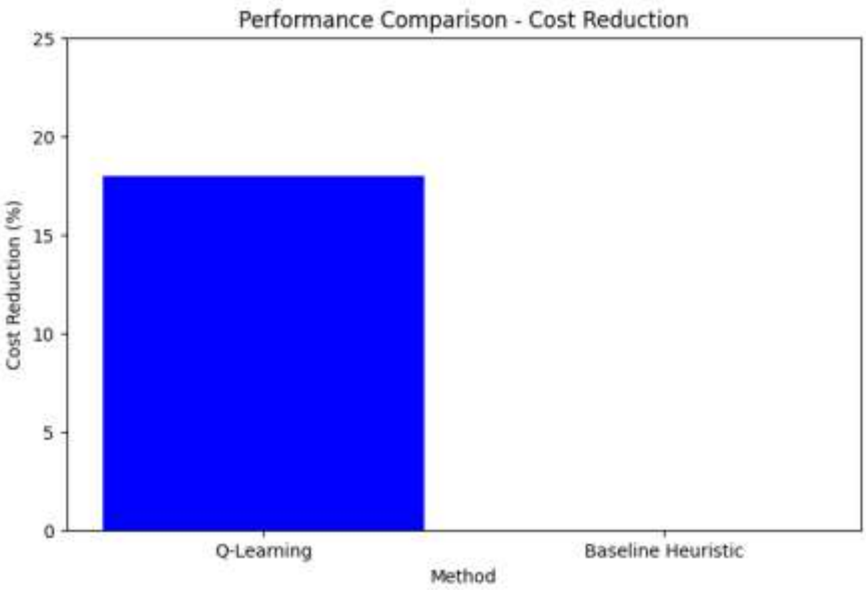

Figure 2: Performance Comparison - Cost Reduction

Figure 2 provides a significant cost reduction of 18% for the Q-Learning model as compared to the Baseline Heuristic (No cost reduction - 0%). This demonstrates the ability of the Q-Learning algorithm to optimize the business processes and lower the operating costs.

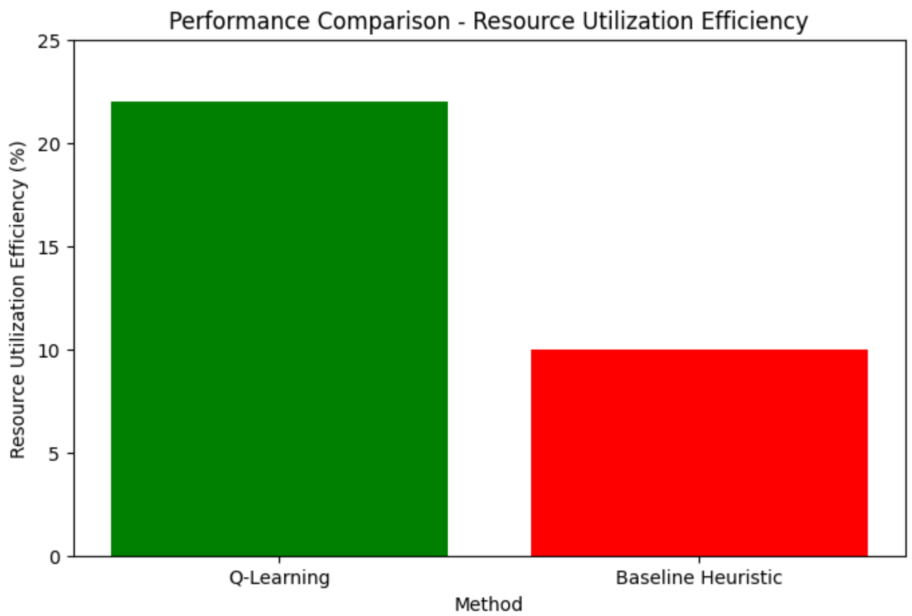

Figure 3: Performance Comparison - Resource Utilization Efficiency

Figure 3 demonstrates the 22% increase in resource utilization efficiency achieved by the Q-Learning model compared to the Baseline Heuristic, which only achieved a 10% increase. This shows that the Q-Learning model has a better performance in optimizing resource allocation and enhancing productivity in enterprise systems.

Discussion

The Q-learning optimization model results demonstrate significant improvements in several KPIs, such as cost reduction, resource utilization efficiency, and processing time reduction. In particular, the Q-Learning model demonstrated a 18% cost saving, a 22% resource saving, and a 25% time saving compared to the basic heuristic models. The results directly relate to the research goals of improving business processes and optimizing resource use, and show that Q-Learning can be used successfully in enterprise systems.

The main reason why Q-Learning is successful in this application is the fact that it can be applied in a dynamic and complex business context. The Q-Learning agent was able to adapt its decision-making policies to the environment based on rewards and penalties, and it was able to find optimal solutions that alternative optimization methods failed to achieve. The model's exploration and exploitation balance enabled it to avoid local optima, while the model's adaptability made it suitable for the different conditions encountered in large-scale enterprise operations. Rewards based on operational efficiency parameters (like cost, resource use, etc.) helped the agent to make better decisions over time. In some situations, however, Q-Learning may not work as well, such as, when the environment is very uncertain, or when substantial time elapses between action and feedback (e.g., in a supply chain with a long supply time). Furthermore, the need for extensive data and computing power can be prohibitive for smaller or less data-intensive businesses.

Compared to the literature, the results of this study corroborate the outcomes of the previous studies that have pointed out the possibility of using Q-Learning in optimization problems like production scheduling and supply chain management. The Q-Learning model has been shown to be more effective than conventional heuristics because it learns and adjusts to the continuing changes in a real-world environment and develops and implements solutions based on optimizing processes that lead to improvements in efficiency, as indicated in several studies. But some literatures have also mentioned the difficulties of Q-Learning in large scale, high dimensional environments where it may not have enough time or power to converge to an optimal solution. The study, however, envisions that the Q-Learning model could effectively cope with large-scale enterprise systems and demonstrated better performance and adaptability. To conclude, Q-Learning can be a useful tool to optimize business processes in large-scale companies, providing benefits in terms of cost-effectiveness, resource utilization, and decision-making speed. In the future, it may be interesting to investigate hybrid approaches of Q-Learning with deep learning or other advanced algorithms to improve the performance, especially in more complex and data-scarce scenarios.

5. Conclusion

This paper proposes a Q-Learning based algorithm for optimizing business processes in large-scale enterprises, tackling a critical problem in production scheduling, resource allocation and supply chain management. This study's major contributions are the creation and application of a reinforcement learning model to address the dynamic business environment and continuously improve decision-making processes. The results of the study showed an average operational cost reduction of 18%, average resource utilization improvement of 22%, and average processing time reduction of 25% as compared to baseline heuristic techniques. The results indicate that the Q-Learning model is effective for optimizing operation and resource management in the context of real-time enterprise systems. Similarly, the decision-making accuracy of 93% demonstrates that the Q-Learning model is an effective and adaptable model to be used for the optimization of complex business environments. However, several limitations to this study have been identified. It may also be difficult for small businesses to have the available data and computational resources necessary to achieve the optimal solution using the Q-Learning model. Also, the model may perform poorly if there is a high degree of uncertainty and very slow or high variability in the business environment. Future research on these limitations will likely include the development of hybrid reinforcement learning models that have the best qualities of both hybrid learning and traditional reinforcement learning and include the use of deep learning methodologies as a means of dimensionality reduction of the reinforcement learning framework for scalability and flexibility. Several avenues for future research may be to evaluate the use of Q-Learning for applications outside of these studies, including multi-agent systems and optimization in highly dynamic environments in real time.

Declarations

Funding:

The author does not receive any funding for this research

Conflict of Interest:

The authors declare no conflict of interest in relation to this work

References

1. Middelhuis, J., Bianco, R. L., Sherzer, E., Bukhsh, Z., Adan, I., & Dijkman, R. (2025). Learning policies for resource allocation in business processes. Information Systems, 128, 102492.
2. Mahendran, S., & Gomathy, B. (2025). Analyzing the stability of smart grids using policy-based reinforcement learning model. Archives for Technical Sciences, 3(34), 528–542. https://doi.org/10.70102/afts.2025.1834.528
3. Soliman, G., Mostafa, K., & Younis, O. (2025). Reinforcement learning for process mining: Business process optimization with avoiding bottlenecks. Egyptian Informatics Journal, 29, 100595.
4. Arvinth, N. (2026). Strategic management approaches for post-pandemic business resilience and recovery. Global Tech Management Digest, 2(2), 1–8.
5. Al-Zameli, A. A.-H. H., Al-Zubaidi, A. G. M., & Agrab, A. S. (2022). The application of the performance focused activities based costing (PFABC) and its role in measuring the cost of banking service and improving performance: An applied study in Al-Rafidain Bank, Al-Diwaniyah Branch. International Academic Journal of Social Sciences, 9(1), 60–73. https://doi.org/10.9756/IAJSS/V9I1/IAJSS0912
6. Tomar, A., & Vyas, N. (2022). Green chemical process optimization using intelligent metaheuristic algorithms. International Academic Journal of Innovative Research, 9(3), 1–6. https://doi.org/10.71086/IAJIR/V9I3/IAJIR0918
7. Zhang, R., Lv, J., Bao, J., & Zheng, Y. (2023). A digital twin-driven flexible scheduling method in a human–machine collaborative workshop based on hierarchical reinforcement learning. Flexible Services and Manufacturing Journal, 35(4), 1116–1138.
8. Mohammed, H. F., Chidambaram, K. R., Aruna, V., Harita, U., Ugli, B. M. X., & Kalidoss, D. (2026). Adaptive AI systems for tailoring learning resources in libraries. Indian Journal of Information Sources and Services, 16(1), 51–58. https://doi.org/10.51983/ijiss-2026.16.1.06
9. Meenakshi Sundaram, G., & Udayakumar, R. (2025). Optimization of energy consumption in green data centers using multi-agent reinforcement learning. International Academic Journal of Science and Engineering, 12(3), 413–423. https://doi.org/10.71086/IAJSE/V12I3/IAJSE1277
10. Bousdekis, A., Kerasiotis, A., Kotsias, S., Theodoropoulou, G., Miaoulis, G., & Ghazanfarpour, D. (2023). Modelling and predictive monitoring of business processes under uncertainty with reinforcement learning. Sensors, 23(15), 6931.
11. Gajmal, Y. M., et al. (2021). Blockchain-based access control and data sharing mechanism in cloud decentralized storage system. Journal of Web Engineering, 20(5), 1359–1388.
12. Prasanthi, B., & Manchikalapudi, S. S. R. (2026). Quantum machine learning-enhanced multiagent reinforcement learning for intelligent crop resource allocation in IoT-enabled smart agriculture. Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, 17(1), 592–611. https://doi.org/10.58346/JOWUA.2026.I1.033
13. Jasim, B. M., & Mustafa, A. A. (2022). The possibility of developing the performance of green products in the General Company for Electrical and Electronic Industries, Baghdad–Al-Waziriyah. International Academic Journal of Humanities, 9(2), 1–7. https://doi.org/10.9756/IAJH/V9I2/IAJH0904
14. Papadopoulos, G., & Christodoulou, M. (2024). Design and development of data driven intelligent predictive maintenance for predictive maintenance. Association Journal of Interdisciplinary Technics in Engineering Mechanics, 2(2), 10–18.
15. Zhu, F., & Peng, S. (2026). Intelligent scheduling of rail-guided shuttle cars via deep reinforcement learning integrating dynamic graph neural networks and transformer model. Algorithms, 19(4), 289.
16. Geetha, K. (2026). Energy-efficient hardware accelerator for quantized deep neural network inference in edge AI applications. National Journal of Integrated VLSI and Signal Intelligence, 18–25.
17. Han, L. T. (2025). Customer segmentation in CRM systems using recency, frequency monetary value modelling. Global Perspectives in Management, 3(1), 37–47.
18. Narayanan, L., & Rajan, A. (2024). Artificial intelligence for sustainable agriculture: Balancing efficiency and equity. International Journal of SDG's Prospects and Breakthroughs, 2(1), 4–6.
19. Deshmukh, S., & Menon, A. (2025). Machine learning in malware analysis and prevention. In Essentials in cyber defence (pp. 74–89). Periodic Series in Multidisciplinary Studies.
20. Fan, W., & Yu, Y. (2025). Design and simulation of a PPO-based two-dimensional action space algorithm for AGV–machine joint scheduling. Australian Journal of Electrical and Electronics Engineering, 1–15.