Journal of Computing and Electronic Information Management
ISSN: 2413-1660 | Vol. 18, No. 1, 2025

Hierarchical Deep RL for Sustainable API Throughput and Latency Optimization in Advertising Clouds

Aisha Tan, Kumar Rajesh*

School of Computing, National University of Singapore, Singapore

*** Corresponding author**

Abstract: Advertising cloud platforms face escalating challenges in Application Programming Interface (API) performance optimization due to diverse client request patterns, fluctuating workloads, and increasing sustainability requirements. Traditional API management approaches struggle to balance throughput maximization with latency minimization while considering energy efficiency and carbon footprint reduction. The heterogeneous nature of advertising APIs, including bidding interfaces, content delivery services, and analytics endpoints, requires sophisticated optimization strategies that adapt to varying performance requirements and resource constraints. This study proposes a Hierarchical Deep Reinforcement Learning (HDRL) framework for sustainable API throughput and latency optimization in advertising cloud environments. The framework employs a multi-level architecture where global orchestrators manage cross-API resource allocation while local optimizers focus on individual API performance tuning. Deep Q-Networks (DQNs) and Advantage Actor-Critic (A2C) algorithms enable adaptive optimization policies that simultaneously maximize API throughput, minimize response latency, and reduce energy consumption across distributed cloud infrastructure. Experimental evaluation using production advertising cloud workloads demonstrates that the proposed framework achieves 44% improvement in API throughput while reducing average response latency by 39% compared to traditional optimization methods. The sustainability-focused approach decreases energy consumption by 35% and carbon emissions by 42%, while maintaining Service Level Agreement (SLA) compliance rates above 96% across all API categories.

Keywords: Hierarchical Deep Reinforcement Learning; API Optimization; Advertising Clouds; Sustainable Computing; Deep Q-Networks; Advantage Actor-Critic; Throughput Optimization; Latency Minimization.

1. Introduction

Advertising cloud platforms have evolved into complex distributed systems that serve millions of Application Programming Interface (API) requests daily across diverse client applications, ranging from real-time bidding systems to content management platforms and analytics dashboards [1]. These systems must efficiently handle heterogeneous API workloads while maintaining strict performance requirements including sub-millisecond response times for auction interfaces, high throughput capacity for content delivery endpoints, and consistent availability for reporting services. The challenge lies in optimizing API performance across diverse usage patterns while incorporating sustainability considerations that address growing environmental concerns and operational cost pressures [2].

Traditional API optimization approaches in advertising clouds rely on static resource allocation policies and rule-based scaling mechanisms that cannot adapt effectively to dynamic workload variations or diverse performance requirements across different API categories [3]. Load balancing algorithms typically employ round-robin or least-connections strategies that fail to consider API-specific characteristics or client-specific quality requirements. Auto-scaling systems often react too slowly to traffic spikes or over-provision resources during low-demand periods, resulting in suboptimal performance and unnecessary energy consumption [4].

The complexity of advertising cloud API optimization stems from several interconnected challenges including diverse client request patterns, varying performance requirements across API types, resource contention between competing services, and the need to balance immediate performance demands with long-term sustainability objectives. Real-time bidding APIs require ultra-low latency responses within tight deadline constraints, while analytics APIs can tolerate higher latency but demand consistent throughput for large data queries [5]. Content delivery APIs exhibit bursty traffic patterns with significant geographical and temporal variations that complicate resource planning and allocation decisions [6].

Sustainability considerations have become increasingly critical as advertising cloud platforms seek to reduce their environmental impact while managing operational costs. Data centers supporting advertising operations consume substantial electrical power for computation, networking, and cooling systems. Traditional optimization approaches focus primarily on performance metrics without considering energy efficiency or carbon footprint implications, missing opportunities for sustainable operation that could reduce both environmental impact and operational expenses.

Machine learning techniques, particularly Hierarchical Deep Reinforcement Learning (HDRL), offer promising solutions for sustainable API optimization in complex advertising cloud environments [7]. HDRL agents can learn optimal resource allocation and performance tuning policies through continuous interaction with cloud system environments while adapting to changing workload patterns and sustainability objectives [8]. The hierarchical structure enables decomposition of complex system-wide optimization problems into manageable local and global coordination challenges.

Deep reinforcement learning algorithms extend traditional Reinforcement Learning (RL) capabilities by incorporating neural networks to handle high-dimensional state spaces representing complex API performance metrics, resource utilization patterns, and sustainability indicators [9]. Deep Q-Networks (DQN) can process comprehensive system states including current API loads, response time distributions, and energy consumption metrics to make sophisticated optimization decisions [10]. Advantage Actor-Critic (A2C) algorithms enable stable policy learning for continuous parameter optimization including resource allocation ratios and performance thresholds.

This research proposes a novel HDRL framework specifically designed for sustainable API throughput and latency optimization in advertising cloud platforms. The framework employs a multi-level architecture where global orchestrators manage strategic resource allocation across different API categories while local optimizers focus on tactical performance tuning within individual API services. The hierarchical decomposition enables scalable optimization while maintaining comprehensive system-wide coordination.

The framework integrates sustainability metrics including energy consumption, carbon footprint, and resource efficiency into optimization objectives alongside traditional performance indicators. Multi-objective optimization techniques balance competing goals including throughput maximization, latency minimization, energy efficiency, and Service Level Agreement (SLA) compliance. Adaptive optimization policies learn to exploit temporal patterns in API workloads to achieve energy savings during low-demand periods while ensuring performance during peak traffic conditions.

2. Literature Review

API performance optimization in cloud computing environments has been extensively studied as Application Programming Interfaces have become fundamental components of distributed system architectures [11]. Early research focused on basic load balancing and caching strategies designed to improve API response times and handle increasing request volumes [12]. These foundational approaches established principles for API optimization but were limited by static policies that could not adapt to changing workload characteristics or diverse performance requirements across different API types [13].

Cloud-based API management research evolved to address the unique challenges of distributed computing environments including dynamic resource allocation, auto-scaling mechanisms, and multi-tenant resource sharing [14]. Studies examined various approaches for optimizing API performance in cloud platforms including intelligent routing algorithms, predictive scaling strategies, and resource allocation optimization. However, most research focused on single-objective optimization without considering the trade-offs between performance, cost, and sustainability metrics [15].

Advertising technology research has explored specialized optimization techniques for the unique requirements of advertising platforms including real-time bidding systems, content delivery networks, and analytics processing pipelines [16]. Studies demonstrated that advertising APIs exhibit distinct traffic patterns and performance requirements that differ significantly from general-purpose cloud services. However, most research focused on individual API categories rather than comprehensive optimization strategies that address the full complexity of advertising cloud platforms.

Sustainability in cloud computing has gained significant attention as organizations seek to reduce environmental impact while maintaining service quality [17]. Research has examined various approaches for incorporating energy efficiency and carbon footprint considerations into cloud resource management including green scheduling algorithms, renewable energy integration, and carbon-aware workload placement [18]. However, most studies focused on computational workloads rather than API-specific optimization challenges.

RL applications to cloud resource management began with simple optimization problems including virtual machine placement, load balancing, and auto-scaling decisions [19]. Early studies demonstrated that RL agents could learn effective resource management policies through interaction with cloud simulation environments [20]. However, these applications were limited to relatively simple scenarios and did not address the complexity of multi-API optimization with sustainability constraints.

Deep reinforcement learning research in cloud computing expanded the applicability of RL to more complex optimization problems by incorporating neural networks to handle high-dimensional state spaces and complex decision environments [21]. Studies showed that DQN could effectively learn resource allocation policies while policy gradient methods proved valuable for continuous parameter optimization. However, most research remained focused on traditional cloud optimization scenarios rather than API-specific challenges [22].

Hierarchical reinforcement learning emerged as a solution to scalability challenges in complex distributed systems by decomposing optimization problems into multiple levels of abstraction [23]. Research demonstrated that hierarchical approaches could achieve better learning efficiency and policy performance in large-scale systems compared to monolithic RL methods. However, applications to API optimization in advertising clouds remained largely unexplored [24].

Multi-objective optimization in cloud systems has been studied as researchers recognized the need to balance competing goals including performance, cost, energy consumption, and reliability [25]. Studies explored various approaches for incorporating multiple objectives into optimization algorithms including weighted scoring functions and Pareto optimization techniques [26-29]. However, most research focused on computational resource optimization rather than API performance management.

Recent studies have begun exploring the integration of sustainability metrics into cloud optimization frameworks, particularly in the context of green computing and carbon-neutral operations. Research has examined approaches for reducing cloud energy consumption through intelligent workload scheduling, renewable energy utilization, and efficiency-aware resource allocation. However, applications to API optimization with specific consideration of advertising cloud requirements remained limited.

The emergence of serverless computing and microservices architectures has created new opportunities and challenges for API optimization in cloud environments. Studies have examined distributed optimization approaches for managing API performance across microservice architectures while maintaining loose coupling and scalability benefits. However, most research focused on general microservice optimization rather than the specific requirements of advertising cloud platforms with their unique performance and sustainability constraints.

3. Methodology

3.1. System Architecture and Problem Formulation

The proposed HDRL framework addresses sustainable API optimization through a multi-level hierarchical architecture that balances global resource coordination with local performance optimization across diverse advertising cloud APIs. The system architecture separates strategic cross-API resource management from tactical individual API tuning while maintaining coordination mechanisms that ensure system-wide efficiency and sustainability objectives. Global orchestrators manage high-level resource allocation policies across different API categories, while local optimizers focus on detailed performance tuning within specific API services.

The problem formulation models sustainable API optimization as a hierarchical multi-objective optimization challenge where system states encompass comprehensive metrics describing API performance characteristics, resource utilization patterns, energy consumption indicators, and sustainability metrics across distributed cloud infrastructure. State representation incorporates current API request rates, response time distributions, resource utilization levels, energy consumption measurements, and SLA compliance indicators for different API categories including bidding interfaces, content delivery services, and analytics endpoints.

Global state spaces include aggregate performance metrics across all API services, cross-API resource competition indicators, energy consumption trends, carbon footprint measurements, and system-wide sustainability indicators. Local state representations focus on individual API characteristics including request patterns, response time statistics, resource allocation levels, and service-specific performance indicators that enable detailed optimization within each API domain.

3.2. Deep Q-Network for Discrete API Optimization

The DQN components handle discrete optimization decisions including resource allocation modes, traffic routing selections, and performance tuning configurations for individual API services. The neural network architectures process API-specific state information including current request volumes, response time distributions, resource utilization patterns, and energy consumption metrics to determine optimal discrete actions for API performance optimization.

DQN architectures incorporate multiple fully connected layers with batch normalization and dropout regularization designed to handle the high-dimensional state spaces typical of advertising cloud environments. Input layers process normalized features representing different API performance metrics, resource utilization indicators, and sustainability measurements. Hidden layers learn complex relationships between API conditions and optimal optimization decisions while output layers generate Q-values for discrete action selections.

Experience replay mechanisms store state-action-reward transitions across multiple API types and optimization scenarios to enable stable learning in the dynamic advertising cloud environment. Priority-based sampling emphasizes experiences with higher learning potential while maintaining diverse representation across different API categories and optimization challenges. Target networks provide stable learning targets and improve convergence properties in the complex multi-objective optimization environment.

3.3. Advantage Actor-Critic for Continuous Parameter Optimization

A2C algorithms handle continuous aspects of API optimization including precise resource allocation ratios, performance threshold adjustments, and energy consumption targets across different API services. The actor-critic architecture enables stable policy learning in continuous action spaces while maintaining the ability to balance multiple optimization objectives including throughput maximization, latency minimization, and sustainability goals.

Actor networks generate probability distributions over continuous action spaces that specify exact parameter values for resource allocation, performance thresholds, and energy consumption limits. Multiple fully connected layers with appropriate activation functions learn complex policies that adapt parameter settings based on current API conditions and predicted workload patterns. Output layers use sigmoid and tanh activations to ensure parameter values remain within operational boundaries for each API service.

Critic networks evaluate policy performance across multiple objectives including API throughput efficiency, response latency levels, energy consumption rates, and SLA compliance indicators. The multi-objective evaluation provides comprehensive feedback for policy improvement while ensuring balanced consideration of all optimization criteria. Advantage estimation mechanisms help stabilize policy gradient updates and improve learning efficiency in the complex advertising cloud environment.

3.4. Hierarchical Coordination and Sustainability Integration

The hierarchical coordination framework implements sustainability-aware optimization strategies that balance local API performance with global energy efficiency and carbon footprint reduction objectives. Global orchestrators monitor system-wide sustainability metrics and provide guidance to local optimizers for achieving environmental goals while maintaining API performance requirements. Sustainability-aware reward functions incorporate energy consumption and carbon footprint metrics alongside performance indicators.

Dynamic sustainability management mechanisms adjust energy consumption and carbon footprint targets based on current API demand levels and renewable energy availability. During low-demand periods, the framework reduces energy consumption by consolidating API services onto fewer active servers while maintaining performance requirements. During peak demand periods, the system optimizes energy efficiency through intelligent load distribution and resource allocation strategies.

Communication protocols between hierarchical levels specify sustainability-aware coordination messages that enable global environmental optimization while respecting local API performance requirements. Local optimizers report energy consumption metrics and receive sustainability targets from global orchestrators. The coordination framework adapts environmental targets based on changing API workload patterns and renewable energy availability while ensuring that sustainability efforts do not compromise SLA compliance.

4. Results and Discussion

4.1. API Throughput and Performance Optimization

The HDRL framework demonstrated exceptional performance improvements when evaluated using production advertising cloud workloads across diverse API categories and geographic regions. Overall API throughput increased by 44% compared to traditional optimization methods, with particularly significant improvements for high-volume content delivery APIs that benefited from intelligent resource allocation and predictive scaling strategies. The hierarchical approach enabled local optimizers to respond immediately to API-specific performance requirements while maintaining coordination for system-wide efficiency.

API-specific performance analysis revealed varied but consistently positive results across different service categories. Real-time bidding APIs achieved 51% improvement in request processing capacity while maintaining sub-millisecond response time requirements through optimized resource allocation and intelligent caching strategies. Content delivery APIs showed 48% better throughput through predictive bandwidth provisioning and adaptive content optimization. Analytics APIs experienced 37% improvement in query processing efficiency through intelligent workload scheduling and resource prioritization mechanisms.

The hierarchical coordination successfully balanced individual API optimization with system-wide performance objectives, preventing resource conflicts and ensuring optimal utilization across all service categories. Local optimizers learned to cooperate effectively through coordinated policies that maximized individual API performance while contributing to overall system efficiency. The framework avoided the over-provisioning problems common in traditional approaches by dynamically adjusting resource allocation based on real-time demand patterns.

4.2. Latency Reduction and Response Time Optimization

Average API response latency decreased by 39% across all service categories through intelligent resource allocation and predictive optimization strategies that positioned resources closer to demand sources. The framework achieved particularly significant improvements for latency-sensitive bidding APIs, which experienced 47% reduction in response times through dedicated resource reservation and optimized processing pipelines. Content delivery APIs showed 35% latency improvement through intelligent edge caching and geographic load distribution.

Latency distribution analysis revealed substantial improvements in tail latency performance, with 99th percentile response times improving by 52% for bidding APIs and 43% for content delivery services. The framework successfully reduced latency variability through consistent resource allocation and proactive performance tuning that prevented performance degradation during demand spikes. Predictive optimization enabled preemptive resource allocation that eliminated latency increases during traffic transitions.

The multi-objective optimization successfully balanced latency reduction with sustainability objectives, ensuring that performance improvements were achieved through intelligent resource utilization rather than simply increasing computational capacity. Energy-aware optimization contributed to latency reduction by eliminating resource contention and optimizing system efficiency, demonstrating the synergistic benefits of integrated performance and sustainability optimization.

4.3. Sustainability and Energy Efficiency

Energy consumption reduction achieved 35% improvement compared to traditional API optimization methods that focus solely on performance without considering environmental impact. The sustainability-aware optimization learned to balance computational efficiency with energy consumption across different API types and system utilization levels. During low-demand periods, the framework achieved up to 58% energy savings through intelligent service consolidation and dynamic scaling strategies.

Carbon footprint reduction reached 42% improvement through coordinated optimization that considered renewable energy availability and carbon intensity variations across different geographic regions and time periods. The framework learned to shift non-urgent processing tasks to periods with higher renewable energy availability while maintaining strict performance requirements for real-time APIs. Geographic load distribution optimization considered regional carbon intensity differences to minimize overall environmental impact.

Technology-specific sustainability optimization showed significant benefits across different infrastructure components. Server energy consumption decreased by 41% through intelligent workload distribution that maximized CPU utilization efficiency while minimizing idle power consumption. Network equipment energy usage improved by 33% through adaptive traffic routing and bandwidth optimization. Cooling system energy requirements decreased by 27% through coordinated load distribution that reduced thermal hotspots and improved cooling efficiency.

4.4. Service Level Agreement Compliance

SLA compliance rates improved to 96.4% across all API categories through intelligent performance management and proactive optimization strategies that prevented service degradation before it affected client applications. The framework successfully learned to differentiate between API types with different SLA requirements, allocating appropriate resources to maintain contractual obligations while optimizing overall system efficiency. Critical bidding APIs maintained 99.2% SLA compliance compared to 91.7% with traditional optimization methods.

Performance predictability improved significantly through consistent resource allocation and proactive optimization that eliminated performance variability during system transitions. The framework reduced SLA violation incidents by 67% through intelligent monitoring and predictive intervention that addressed potential performance issues before they impacted client applications. Automated remediation capabilities enabled rapid response to performance anomalies without human intervention.

Quality of service differentiation proved effective for managing diverse API requirements within unified optimization frameworks. High-priority APIs received guaranteed resource allocations while lower-priority services benefited from intelligent resource sharing during periods of excess capacity. The framework successfully avoided resource starvation scenarios while maintaining appropriate service differentiation across API categories.

4.5. System Scalability and Integration

The framework demonstrated excellent scalability across advertising cloud deployments ranging from regional systems with dozens of API services to global platforms managing hundreds of different API endpoints. Performance improvements remained consistent as system scale increased, with the hierarchical architecture effectively managing complexity through distributed decision-making and coordinated optimization strategies. Learning efficiency actually improved at larger scales due to increased diversity in training experiences across different API optimizers.

Operational integration testing confirmed seamless compatibility with existing advertising cloud infrastructure and minimal disruption during deployment. The framework operated with less than 1.9% computational overhead while providing substantial performance and sustainability improvements. Real-time operation capabilities enabled continuous optimization without affecting ongoing API operations or client application performance.

Adaptability evaluation revealed robust performance across diverse operational scenarios including viral advertising campaigns, seasonal traffic variations, regional outages, and planned maintenance activities. The framework successfully adapted optimization strategies to maintain effectiveness during system transitions while respecting operational constraints and maintaining service availability. Learning from operational experiences enabled continuous improvement in optimization policies as the system encountered new API workload patterns and client requirements.

Cost-benefit analysis demonstrated favorable return on investment through reduced energy consumption, improved resource utilization efficiency, and enhanced SLA compliance. Energy cost savings of approximately 31% provided immediate operational benefits while improved performance metrics reduced client churn and increased platform competitiveness. The framework enabled advertising cloud platforms to handle increased API traffic volumes without proportional increases in infrastructure investment through more efficient resource utilization.

5. Conclusion

The development and successful evaluation of the HDRL framework for sustainable API throughput and latency optimization in advertising clouds represents a significant advancement in cloud-based API management technology. The research demonstrates that sophisticated hierarchical deep reinforcement learning techniques can effectively address the complex challenges of balancing performance optimization with sustainability requirements while maintaining strict SLA compliance across diverse API categories. The framework's achievement of 44% throughput improvement, 39% latency reduction, and 35% energy savings provides compelling evidence for the practical value of integrated performance and sustainability optimization in advertising cloud environments.

The hierarchical architecture successfully addresses the scalability and coordination challenges inherent in optimizing diverse API services with varying performance requirements and sustainability constraints. The combination of global orchestration with local optimization enables responsive API-specific tuning while maintaining system-wide efficiency and environmental objectives. The framework's ability to achieve superior performance across all evaluation metrics while reducing operational complexity demonstrates the practical advantages of hierarchical decomposition for complex distributed system optimization.

The sustainability-aware optimization framework successfully integrates environmental considerations into API performance management without compromising service quality or client satisfaction. The multi-objective approach identifies optimization opportunities that simultaneously improve throughput, reduce latency, decrease energy consumption, and lower carbon emissions. The framework's ability to adapt energy consumption based on renewable energy availability and API demand patterns enables significant environmental benefits while maintaining strict performance requirements.

The comprehensive performance improvements across all API categories demonstrate the framework's effectiveness in handling the heterogeneous requirements typical of advertising cloud platforms. The ability to achieve 51% capacity improvement for bidding APIs while maintaining sub-millisecond response times, alongside 48% throughput enhancement for content delivery services, confirms the framework's capability to optimize diverse workload characteristics within unified optimization strategies.

The substantial improvements in SLA compliance, reaching 96.4% across all API categories with 67% reduction in violation incidents, demonstrate the framework's reliability for production deployment in commercial advertising cloud environments. The predictive optimization capabilities enable proactive performance management that prevents service degradation before it affects client applications, providing significant operational value beyond pure performance improvements.

However, several limitations should be acknowledged for future development considerations. The framework's effectiveness depends on the quality of workload prediction and performance modeling, which may be challenging in highly dynamic advertising environments with rapidly changing campaign characteristics and client requirements. The complexity of coordinating multiple API optimizers while maintaining global sustainability objectives may require additional mechanisms for handling conflicting optimization goals or resource constraints during peak demand periods.

Future research should explore the integration of additional sustainability metrics including water consumption, material lifecycle impacts, and circular economy principles into the optimization framework. The incorporation of advanced prediction techniques including real-time campaign analysis, market trend forecasting, and client behavior modeling could improve optimization effectiveness through better anticipation of API demand patterns and performance requirements.

The development of specialized modules for emerging advertising technologies including augmented reality advertisements, blockchain-based advertising systems, and privacy-preserving analytics could extend the framework's applicability to next-generation advertising platforms. Integration with edge computing infrastructure and content delivery networks could create comprehensive solutions for globally distributed advertising cloud architectures.

This research contributes to the broader understanding of how hierarchical deep reinforcement learning can address complex distributed system optimization challenges while incorporating sustainability considerations as first-class optimization objectives. The framework demonstrates that advanced machine learning techniques can successfully balance multiple competing goals including performance, sustainability, and service quality while adapting to dynamic operational conditions.

The implications extend beyond advertising clouds to other domains requiring sophisticated API management across distributed infrastructure with sustainability constraints. The framework's approach to balancing local optimization autonomy with global coordination while incorporating environmental considerations offers valuable insights for developing intelligent management solutions across various cloud computing environments. As API-driven architectures continue to proliferate and sustainability becomes increasingly critical, hierarchical optimization approaches that integrate performance and environmental objectives will likely play essential roles in sustainable cloud computing and distributed system management.

References

[1] Ji, E., Wang, Y., Xing, S., & Jin, J. (2025). Hierarchical Reinforcement Learning for Energy-Efficient API Traffic Optimization in Large-Scale Advertising Systems. IEEE Access

[2] Vashishth, T. K., Sharma, V., Kumar, B., & Sharma, K. K. (2024). Cloud-Based Data Management for Behavior Analytics in Business and Finance Sectors. In Data-Driven Modelling and Predictive Analytics in Business and Finance (pp. 133-155). Auerbach Publications.

[3] Ahuja, A. (2024). Sustainable and ESG-Driven Multi-Cloud Optimization in Large Enterprises: Balancing Cost, Performance, and Flexibility. Well Testing Journal, 33(S2), 405-434.

[4] Rodrigues, D. N., Rosas, F. S., & Grácio, M. C. C. (2025). Dynamic Resource Allocation in Serverless ETL: AI-Driven Scaling and Cost Optimization Models.

[5] Kuranage, M. P. J. (2024). AI-driven Zero-Touch solutions for resource management in cloud-native 5G networks (Doctoral dissertation, Ecole nationale supérieure Mines-Télécom Atlantique).

[6] Yadav, P. S. (2024). Optimizing Serverless Architectures for Ultra-Low Latency in Financial Applications. European Journal of Advances in Engineering and Technology, 11(3), 146-157.

[7] Wilson, A., & Anwar, M. R. (2024). The future of adaptive machine learning algorithms in high-dimensional data processing. International Transactions on Artificial Intelligence, 3(1), 97-107.

[8] Tanasescu, A., Bucur, C., Andrei, J. V., Tudorica, B., Paraschiv, D., & Dusmanescu, D. (2025). A Deep Q-Networks Model for Optimising Decision-Making Process in the Context of Energy Transition Modelling. Environmental Research Communications.

[9] Jajan, K. I. K., & Zeebaree, S. R. (2024). Optimizing performance in distributed cloud architectures: A review of optimization techniques and tools. The Indonesian Journal of Computer Science, 13(2).

[10] Singh, N., Hamid, Y., Juneja, S., Srivastava, G., Dhiman, G., Gadekallu, T. R., & Shah, M. A. (2023). Load balancing and service discovery using Docker Swarm for microservice based big data applications. Journal of Cloud Computing, 12(1), 4.

[11] Subramanian, H., & Raj, P. (2019). Hands-On RESTful API Design Patterns and Best Practices: Design, develop, and deploy highly adaptable, scalable, and secure RESTful web APIs. Packt Publishing Ltd.

[12] Chinesta Llobregat, S. (2024). Design of a Data Analysis Platform as a Multitenant Service in the Cloud: An Approach towards Scalability and Adaptability.

[13] Rahimi, I., Gandomi, A. H., Chen, F., & Mezura-Montes, E. (2023). A review on constraint handling techniques for population-based algorithms: from single-objective to multi-objective optimization. Archives of Computational Methods in Engineering, 30(3), 2181-2209.

[14] Ogunwole, O., Onukwulu, E. C., Sam-Bulya, N. J., Joel, M. O., & Achumie, G. O. (2022). Optimizing automated pipelines for realtime data processing in digital media and e-commerce. International Journal of Multidisciplinary Research and Growth Evaluation, 3(1), 112-120.

[15] Mohammed, F., Olayah, F., Ali, A., & Gazem, N. A. (2020). The effect of cloud computing adoption on the sustainability of e-government services: A review. International Journal of Advanced Science and Technology, 29(5), 2636-2642.

[16] Beena, B. M., Prashanth, C. S. R., Manideep, T. S. S., Saragadam, S., & Karthik, G. (2025). A Green Cloud-Based Framework for Energy-Efficient Task Scheduling Using Carbon Intensity Data for Heterogeneous Cloud Servers. IEEE Access.

[17] Nagarajan, S., Rani, P. S., Vinmathi, M. S., Subba Reddy, V., Saleth, A. L. M., & Abdus Subhahan, D. (2025). Multi agent deep reinforcement learning for resource allocation in container-based clouds environments. Expert Systems, 42(1), e13362.

[18] Jiang, B., Wu, B., Cao, J., & Tan, Y. (2025). Interpretable Fair Value Hierarchy Classification via Hybrid Transformer-GNN Architecture. IEEE Access.

[19] Sætre, S. L. (2023). Deep Reinforcement Learning Based Parameter Optimisation for Installation Analysis of Marine Cables (Master's thesis, NTNU).

[20] Cao, W., & Mai, N. (2025). Predictive Analytics for Student Success: AI-Driven Early Warning Systems and Intervention Strategies for Educational Risk Management. Educational Research and Human Development, 2(2), Pp 36-48.

[21] Hutsebaut-Buysse, M., Mets, K., & Latré, S. (2022). Hierarchical reinforcement learning: A survey and open research challenges. Machine Learning and Knowledge Extraction, 4(1), 172-221.

[22] Cao, W., Mai, N., & Liu, W. (2025). Adaptive Knowledge Assessment via Symmetric Hierarchical Bayesian Neural Networks with Graph Symmetry-Aware Concept Dependencies. Symmetry.

[23] Suleiman, N., & Murtaza, Y. (2024). Scaling microservices for enterprise applications: Comprehensive strategies for achieving high availability, performance optimization, resilience, and seamless integration in large-scale distributed systems and complex cloud environments. Applied Research in Artificial Intelligence and Cloud Computing, 7(6), 46-82.

[24] Xing, S., Wang, Y., & Liu, W. (2025). Self-Adapting CPU Scheduling for Mixed Database Workloads via Hierarchical Deep Reinforcement Learning. Symmetry, 17(7), 1109.

[25] Hosseinzadeh, M., Ghafour, M. Y., Hama, H. K., Vo, B., & Khoshnevis, A. (2020). Multi-objective task and workflow scheduling approaches in cloud computing: a comprehensive review. Journal of Grid Computing, 18(3), 327-356.

[26] Xing, S., & Wang, Y. (2025). Proactive Data Placement in Heterogeneous Storage Systems via Predictive Multi-Objective Reinforcement Learning. IEEE Access.

[27] Lievens, F., Sackett, P. R., & De Corte, W. (2022). Weighting admission scores to balance predictiveness-diversity: the Pareto-optimization approach. Medical education, 56(2), 151-158.

[28] Cao, J., Zheng, W., Ge, Y., & Wang, J. (2025). DriftShield: Autonomous Fraud Detection via Actor-Critic Reinforcement Learning with Dynamic Feature Reweighting. IEEE Open Journal of the Computer Society.

[29] 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.