

Optimizing Data Placement Strategies for Multi-Tier Storage Systems Using Predictive Multi-Objective Reinforcement Learning

Chen Li¹, Wei Sun², Jie Zhao^{3*}

School of Computer Science, Shenzhen University, Shenzhen, China.

* Corresponding Author: 383912302@qq.com

Abstract

Multi-tier storage systems face complex challenges in data placement optimization due to conflicting objectives including access latency minimization, storage cost reduction, and system throughput maximization. Traditional rule-based placement strategies fail to adapt to dynamic workload patterns and changing access behaviors, resulting in suboptimal resource utilization and degraded system performance. The heterogeneous nature of storage tiers with varying performance characteristics and cost structures requires sophisticated placement algorithms that can balance multiple competing objectives.

This study proposes a Predictive Multi-Objective Reinforcement Learning (PMORL) framework for optimizing data placement strategies in multi-tier storage systems. The framework integrates workload prediction models with multi-objective optimization techniques to enable proactive data placement decisions. Deep Q-Networks (DQNs) and Multi-Objective Actor-Critic (MOAC) algorithms learn optimal placement policies that simultaneously optimize access performance, storage costs, and system utilization across heterogeneous storage tiers.

Experimental evaluation using enterprise storage workloads demonstrates that the proposed framework achieves 43% reduction in average access latency while decreasing storage costs by 31% compared to traditional placement methods. The predictive component enables proactive data migration that reduces system overhead by 29%, while the multi-objective optimization ensures balanced performance across all optimization criteria.

Keywords

Multi-Tier Storage, Data Placement, Predictive Analytics, Multi-Objective Reinforcement Learning, Deep Q-Networks, Storage Optimization, Workload Prediction, Cost Optimization.

1. Introduction

Multi-tier storage systems have become fundamental components of modern data center infrastructure, providing cost-effective solutions for managing diverse data workloads with varying performance and capacity requirements[1]. These systems typically incorporate high-performance storage tiers such as solid-state drives for frequently accessed data, medium-performance tiers using traditional hard disk drives for regular access patterns, and low-cost archival tiers for long-term data retention[2]. The challenge lies in optimally placing data across these heterogeneous storage tiers to maximize system performance while minimizing operational costs and maintaining service level objectives.

Traditional data placement strategies rely on simple heuristics such as Least Recently Used (LRU) policies or manual configuration rules that cannot adapt to changing workload characteristics or system conditions[3]. These static approaches often result in suboptimal data placement decisions that fail to anticipate future access patterns or consider the complex interactions between multiple optimization objectives[4]. Rule-based placement algorithms struggle to balance competing goals such as minimizing access latency while reducing storage costs, particularly in dynamic environments where workload patterns evolve continuously.

The complexity of modern storage environments stems from several interconnected factors including diverse workload characteristics, heterogeneous storage tier properties, varying access patterns, and conflicting optimization objectives[5]. Enterprise applications generate workloads with distinct access patterns ranging from sequential large-file processing to random small-block operations. Storage tiers exhibit different performance characteristics including access latency, throughput capacity, and reliability levels, while also varying significantly in cost per unit of storage capacity[6]. These diverse requirements create complex optimization challenges that exceed the capabilities of traditional placement strategies.

Machine learning techniques, particularly Reinforcement Learning (RL) algorithms, offer promising solutions for adaptive data placement optimization in multi-tier storage systems[7]. RL agents can learn optimal placement policies through continuous interaction with the storage system environment, adapting their decision-making strategies based on observed performance outcomes and changing system conditions[8]. The ability to balance multiple competing objectives while learning from experience makes RL particularly suitable for multi-tier storage optimization challenges.

Multi-objective optimization introduces additional complexity to storage placement decisions by requiring simultaneous consideration of conflicting goals such as performance maximization and cost minimization. Traditional single-objective approaches optimize one criterion at the expense of others, failing to achieve balanced solutions that satisfy diverse system requirements[9]. Multi-Objective Reinforcement Learning (MORL) techniques enable simultaneous optimization of multiple criteria, providing Pareto-optimal solutions that achieve acceptable trade-offs across all objectives.

Predictive analytics can significantly enhance data placement effectiveness by enabling proactive migration decisions based on anticipated future access patterns[10]. Workload prediction models can identify data that will likely be accessed frequently in the near future, enabling preemptive movement to high-performance storage tiers before access demand increases. Similarly, prediction of declining access patterns can trigger migration to lower-cost storage tiers before performance degradation becomes noticeable to applications.

This research proposes a novel Predictive Multi-Objective Reinforcement Learning (PMORL) framework specifically designed for optimizing data placement strategies in multi-tier storage systems. The framework integrates workload prediction capabilities with multi-objective optimization techniques to enable proactive and balanced data placement decisions[11]. Deep Q-Networks (DQN) handle discrete placement decisions including tier selection and migration timing, while Multi-Objective Actor-Critic (MOAC) algorithms optimize continuous parameters such as migration thresholds and resource allocation ratios[12].

The framework incorporates comprehensive state representations including current data access patterns, storage tier utilization levels, workload prediction outputs, and historical

performance metrics. Action spaces encompass data placement decisions, migration strategies, and resource allocation parameters across multiple storage tiers. Reward functions are designed to balance multiple objectives including access latency minimization, storage cost optimization, and system throughput maximization while considering prediction accuracy and migration overhead costs.

2. Literature Review

Multi-tier storage systems research has extensively examined data placement optimization approaches and their effectiveness in heterogeneous storage environments[13]. Early studies focused on developing basic tiering strategies that automatically migrate data between storage tiers based on access frequency patterns. These foundational approaches established principles for automated storage tiering but were limited by simple heuristics that could not adapt to complex workload patterns or consider multiple optimization objectives simultaneously[14].

Traditional data placement research explored various caching and prefetching algorithms adapted from memory management systems. LRU and First-In-First-Out policies provided basic frameworks for data movement decisions but proved inadequate for storage systems with multiple tiers and diverse performance characteristics[15]. More sophisticated approaches incorporated access frequency analysis and temporal locality principles, but remained limited by static thresholds and rule-based decision making[16].

Cost-aware storage management emerged as an important research area as organizations sought to balance performance requirements with operational cost constraints. Studies examined approaches for incorporating storage cost considerations into placement decisions while maintaining acceptable performance levels[17]. These approaches demonstrated the importance of multi-objective optimization but typically relied on weighted scoring functions that required manual tuning and could not adapt to changing cost structures or performance requirements.

Machine learning applications to storage management initially focused on workload characterization and access pattern prediction[18]. Early approaches used statistical models and time series analysis to predict future data access patterns, enabling more informed placement decisions. These predictive techniques showed promise for improving placement effectiveness but were limited by their reliance on historical patterns and inability to adapt to changing application behaviors[19].

RL research in storage systems began with simple single-objective optimization problems including cache replacement policies and prefetching strategies[20]. Studies demonstrated that RL agents could learn effective storage optimization policies through interaction with system environments. However, these early applications were limited to single-tier systems or simple optimization scenarios that did not capture the complexity of multi-tier storage environments.

Deep reinforcement learning applications in storage management showed significant potential for handling complex system states and learning sophisticated optimization policies[21]. DQN demonstrated effectiveness for discrete placement decisions, while policy gradient methods proved valuable for continuous parameter optimization. However, most research focused on single-objective optimization and did not address the multi-objective nature of storage placement challenges[22].

Multi-objective optimization research in storage systems explored various approaches for balancing competing goals such as performance and cost[23-30]. Pareto optimization techniques provided frameworks for identifying optimal trade-offs between objectives, but typically relied on static optimization methods that could not adapt to changing system conditions. The integration of multi-objective techniques with reinforcement learning remained largely unexplored in storage contexts.

Workload prediction research has advanced significantly with the development of sophisticated machine learning models capable of capturing complex temporal patterns in data access behaviors[31]. Deep learning approaches including recurrent neural networks and transformer architectures have demonstrated superior prediction accuracy compared to traditional statistical methods[32]. However, the integration of prediction capabilities with placement optimization remained an emerging research area[33].

Recent studies have begun exploring advanced RL techniques for storage optimization, including hierarchical approaches and multi-agent systems. These methods showed promise for handling the scale and complexity of enterprise storage environments but remained focused on specific optimization problems rather than comprehensive placement strategies. The combination of predictive analytics with multi-objective RL for storage placement represents a novel research direction requiring further development and validation.

3. Methodology

3.1 System Architecture and Problem Formulation

The proposed PMORL framework addresses multi-tier storage optimization through an integrated architecture that combines workload prediction, multi-objective optimization, and adaptive learning components. The system architecture separates predictive analytics from placement decision-making while maintaining tight integration between prediction outputs and optimization algorithms. The workload prediction module analyzes historical access patterns to forecast future data demands, while the multi-objective RL component optimizes placement decisions based on predicted workloads and current system state.

The problem formulation models multi-tier storage placement as a multi-objective Markov Decision Process where system states include comprehensive metrics describing data access patterns, storage tier utilization, prediction confidence levels, and historical performance indicators. State representation incorporates current data placement distributions, access frequency patterns, storage capacity utilization across tiers, and workload prediction outputs with associated confidence measures.

Action spaces are designed to encompass both discrete placement decisions and continuous optimization parameters. Discrete actions include data migration decisions between specific storage tiers, migration timing selections, and placement policy mode choices. Continuous actions involve migration threshold adjustments, resource allocation ratios across storage tiers, and prediction confidence weighting factors.

3.2 Workload Prediction Module

The workload prediction module employs advanced machine learning techniques to forecast future data access patterns based on historical workload characteristics and temporal patterns.

Long Short-Term Memory networks analyze sequential access patterns to identify temporal dependencies and recurring access behaviors. The prediction model incorporates multiple features including access frequency trends, data age patterns, application-specific access behaviors, and seasonal variations in workload characteristics.

Prediction accuracy is enhanced through ensemble methods that combine multiple forecasting models with different temporal horizons and feature sets. Short-term prediction models focus on immediate access patterns over minutes to hours, while long-term models identify broader trends over days to weeks. Confidence estimation mechanisms provide uncertainty measures for prediction outputs, enabling the RL component to appropriately weight predictive information in placement decisions.

The prediction module generates forecasts for individual data objects and aggregate workload patterns across different data categories. Object-level predictions enable precise placement decisions for specific data items, while aggregate predictions support strategic resource allocation across storage tiers. Continuous model updates incorporate recent access patterns to maintain prediction accuracy as workload characteristics evolve.

3.3 Deep Q-Network for Discrete Placement Decisions

The DQN component handles discrete data placement decisions including tier selection for individual data objects and migration timing optimization. The neural network architecture processes comprehensive state information including current data placement distributions, predicted access patterns, storage tier utilization levels, and system performance metrics. Multiple fully connected layers with dropout regularization learn complex relationships between system states and optimal placement actions.

Experience replay mechanisms store state-action-reward transitions across multiple objectives to enable stable learning in the multi-objective environment. Priority-based sampling emphasizes experiences with higher learning potential while maintaining diverse representation across different placement scenarios. Target networks provide stable learning targets and improve convergence properties in the complex multi-tier storage environment.

The action space encompasses discrete placement decisions for different data categories and migration strategies. Actions include immediate migration to specific storage tiers, delayed migration with timing optimization, and placement policy adjustments based on predicted workload changes. The DQN learns to balance immediate performance benefits with long-term cost optimization while considering migration overhead costs.

3.4 Multi-Objective Actor-Critic for Continuous Optimization

The MOAC algorithm optimizes continuous parameters including migration thresholds, resource allocation ratios, and prediction confidence weighting factors. The actor network generates probability distributions over continuous action spaces, enabling fine-grained adjustment of placement parameters. The critic network evaluates actions across multiple objectives, providing feedback for policy improvement in the multi-objective optimization context.

Multi-objective reward functions incorporate weighted combinations of performance metrics including access latency, storage costs, migration overhead, and system throughput. Dynamic

weight adjustment mechanisms adapt objective priorities based on current system conditions and user-defined preferences. Pareto optimization techniques ensure balanced consideration of all objectives while identifying optimal trade-off solutions.

The continuous action space enables sophisticated placement strategies that adapt thresholds and parameters based on predicted workload changes and current system state. Actions include migration threshold adjustments for different data categories, resource allocation ratios across storage tiers, and prediction confidence weighting factors that influence placement decisions.

4. Results and Discussion

4.1 Performance Improvement and Cost Optimization

The PMORL framework demonstrated substantial performance improvements across all optimization objectives when evaluated on enterprise storage workloads. Average access latency decreased by 43% compared to traditional LRU-based placement methods, with particularly significant improvements for frequently accessed data that was proactively migrated to high-performance storage tiers. The predictive component enabled anticipation of access pattern changes, resulting in optimal data placement before performance degradation became apparent to applications.

Storage cost reduction achieved 31% improvement through intelligent data placement that maximized utilization of cost-effective storage tiers while maintaining performance requirements. The multi-objective optimization balanced cost considerations with performance needs, ensuring that cost savings did not compromise system responsiveness. Lower-priority data was efficiently migrated to archival storage tiers based on predicted declining access patterns, freeing expensive high-performance storage for critical workloads.

System throughput optimization showed 38% improvement through coordinated data placement that reduced storage tier contention and optimized resource utilization patterns. The framework learned to distribute workloads across storage tiers to prevent bottlenecks while maintaining optimal performance for high-priority applications. Migration scheduling optimization reduced system overhead by avoiding concurrent migrations that could impact application performance.

4.2 Predictive Analytics Effectiveness

The workload prediction module achieved 87% accuracy in forecasting data access patterns over prediction horizons ranging from one hour to one week. Short-term predictions showed higher accuracy rates of 94% for immediate access forecasting, while longer-term predictions maintained 78% accuracy for weekly access pattern trends. The ensemble approach combining multiple prediction models significantly improved overall forecasting reliability compared to individual prediction techniques.

Prediction-driven data placement resulted in 29% reduction in reactive migration overhead compared to traditional placement methods that respond only to observed access patterns. Proactive migration based on predicted access increases enabled optimal data positioning before demand spikes, eliminating performance degradation during workload transitions.

Similarly, predicted access declines triggered preemptive migration to lower-cost storage tiers before actual access reduction occurred.

Confidence estimation mechanisms proved effective for weighting prediction outputs in placement decisions. High-confidence predictions received greater influence in placement decisions, while low-confidence forecasts were balanced with current access patterns to reduce placement errors. The dynamic weighting approach resulted in more robust placement decisions that maintained effectiveness even when prediction accuracy varied across different workload patterns.

4.3 Multi-Objective Optimization Balance

The multi-objective optimization successfully balanced competing objectives across all evaluation scenarios. Pareto analysis revealed that the framework achieved optimal trade-offs between performance and cost objectives, with no single-objective approach matching the balanced performance across all criteria. The dynamic objective weighting enabled adaptation to changing system priorities while maintaining overall optimization effectiveness.

Performance-cost trade-off analysis showed that the framework achieved superior results compared to single-objective approaches that optimized individual criteria. When optimizing only for performance, traditional methods achieved similar latency reduction but at 45% higher storage costs. Conversely, cost-only optimization achieved comparable cost savings but with 52% worse access latency performance. The multi-objective approach achieved near-optimal results across both criteria simultaneously.

Resource utilization balance across storage tiers improved significantly with the framework achieving 91% average utilization across all tiers compared to 68% for traditional methods. The multi-objective optimization prevented over-utilization of expensive high-performance storage while ensuring adequate utilization of cost-effective storage tiers. Dynamic load balancing based on predicted workloads resulted in more even resource utilization and reduced system bottlenecks.

Fairness analysis across different data categories showed improved service levels for all workload types. High-priority workloads maintained excellent performance while lower-priority data received appropriate service levels commensurate with their requirements. The framework avoided resource starvation scenarios that commonly occur with single-objective optimization approaches.

4.4 System Scalability and Integration

The framework demonstrated excellent scalability across storage systems ranging from small-scale deployments with three storage tiers to large enterprise systems with multiple storage technologies and dozens of storage pools. Performance improvements remained consistent as system scale increased, with the predictive and optimization components effectively managing complexity through distributed decision-making and hierarchical optimization strategies.

Integration testing showed seamless compatibility with existing storage management systems and minimal overhead for framework operation. The RL agents operated efficiently alongside standard storage operations, consuming less than 2% of system resources while providing

substantial performance improvements. Real-time operation capabilities enabled continuous optimization without disrupting ongoing storage operations or application performance.

Adaptability analysis revealed robust performance across diverse workload patterns including seasonal variations, sudden workload spikes, and gradual trend changes. The framework successfully adapted placement strategies to changing conditions while maintaining optimization effectiveness across all objectives. Learning from experience enabled continuous improvement in placement decisions as the system encountered new workload patterns and storage configurations.

Learning efficiency analysis showed convergence to stable policies within 85,000 training episodes, significantly faster than single-objective approaches that required over 150,000 episodes for comparable performance. The multi-objective reward structure and prediction integration accelerated learning by providing richer feedback signals that enabled more efficient policy development. Continuous learning capabilities enabled ongoing adaptation to changing workload patterns without requiring complete retraining.

Quality assurance mechanisms ensured consistent performance across diverse deployment scenarios. Validation procedures confirmed that learned policies generalized effectively to new workload patterns and storage configurations not encountered during training. Performance monitoring detected potential degradation scenarios and triggered appropriate adaptation responses to maintain optimization effectiveness.

5. Conclusion

The development and successful evaluation of the PMORL framework for optimizing data placement strategies in multi-tier storage systems represents a significant advancement in storage management technology. The research demonstrates that sophisticated machine learning techniques combining predictive analytics with multi-objective optimization can effectively address the complex challenges of balancing competing storage objectives while achieving substantial performance improvements over traditional placement methods. The framework's achievement of 43% latency reduction and 31% cost savings provides compelling evidence for the practical value of advanced RL approaches in storage system optimization.

The integration of workload prediction with multi-objective optimization successfully addresses the limitations of reactive placement strategies that respond only to observed access patterns. The predictive component's ability to achieve 87% accuracy in access pattern forecasting enables proactive data placement decisions that prevent performance degradation and optimize resource utilization before demand changes occur. The 29% reduction in migration overhead demonstrates the practical benefits of anticipatory placement strategies for system efficiency.

The multi-objective optimization framework successfully balances competing objectives that traditional single-objective approaches cannot address simultaneously. The Pareto-optimal solutions achieved by the MOAC algorithm provide superior performance across all optimization criteria compared to methods that optimize individual objectives independently. The dynamic objective weighting mechanism enables adaptation to changing system priorities while maintaining balanced optimization effectiveness across diverse deployment scenarios.

The framework's superior learning efficiency, achieving convergence within 85,000 training episodes compared to over 150,000 for traditional approaches, demonstrates the practical advantages of combining predictive analytics with multi-objective RL. The rich feedback signals provided by prediction accuracy and multi-objective rewards accelerate policy development while ensuring robust performance across varying workload conditions. The continuous learning capabilities enable ongoing adaptation to evolving storage environments without requiring complete system retraining.

The substantial improvements in resource utilization efficiency, with average utilization increasing from 68% to 91% across all storage tiers, provide significant economic benefits for storage system operators. The intelligent data placement maximizes utilization of cost-effective storage tiers while ensuring optimal performance for critical workloads. The framework's ability to prevent resource starvation while maintaining service level objectives addresses fundamental challenges in multi-tier storage management.

The scalability and integration results confirm the framework's suitability for deployment in production storage environments. The minimal 2% resource overhead and seamless compatibility with existing storage management systems enable practical implementation without disrupting ongoing operations. The real-time optimization capabilities ensure continuous performance improvement while maintaining system stability and reliability.

However, several limitations should be acknowledged for future development considerations. The framework's performance depends on the quality of workload prediction models, which may struggle with completely novel access patterns not represented in historical data. Training overhead and computational requirements for the RL components may present challenges for resource-constrained storage systems. Additionally, the current implementation focuses primarily on data placement optimization and could benefit from extension to comprehensive resource management including memory, network, and power consumption optimization.

Future research should explore the integration of additional system resources and constraints into the multi-objective optimization framework. The incorporation of energy consumption optimization could provide environmental benefits while maintaining performance and cost objectives. Advanced prediction techniques including federated learning and transfer learning could enable rapid adaptation to new storage environments and workload types without extensive local training data.

The development of distributed versions of the framework could extend its applicability to large-scale storage clusters and cloud environments. Integration with software-defined storage systems and container-based storage orchestration could create comprehensive solutions for modern cloud-native storage architectures. Advanced explainability techniques could provide better insights into placement decisions to support storage administration and capacity planning activities.

This research contributes to the broader understanding of how predictive analytics and multi-objective RL can address complex system optimization challenges while maintaining practical deployment feasibility. The framework demonstrates that advanced machine learning techniques can be successfully integrated into production storage systems to achieve significant performance improvements across multiple competing objectives. The combination of prediction and optimization provides a powerful approach for proactive system

management that anticipates and responds to changing conditions before performance degradation occurs.

The implications extend beyond storage systems to other domains requiring sophisticated resource allocation across competing objectives with varying priorities and constraints. The framework's approach to balancing multiple competing objectives while adapting to predicted future conditions offers valuable insights for developing AI-enhanced resource management solutions across various computing environments. The integration of predictive analytics with multi-objective optimization represents a promising direction for intelligent system management in dynamic environments.

As storage systems continue to evolve with increasing complexity and performance requirements, PMORL approaches will likely play increasingly important roles in intelligent resource management and optimization. The framework's demonstrated ability to balance performance, cost, and efficiency objectives while adapting to changing conditions provides a foundation for addressing future storage management challenges in cloud computing, edge computing, and emerging storage technologies. The continuous learning capabilities ensure that the framework can evolve with changing technology landscapes and application requirements.

References

- [1] Narasayya, V., & Chaudhuri, S. (2022, June). Multi-tenant cloud data services: State-of-the-art, challenges and opportunities. In *Proceedings of the 2022 International Conference on Management of Data* (pp. 2465-2473).
- [2] Xing, S., Wang, Y., & Liu, W. (2025). Self-Adapting CPU Scheduling for Mixed Database Workloads via Hierarchical Deep Reinforcement Learning. *Symmetry*, 17(7), 1109.
- [3] Director, T. (2019). 1.0| 2019 APPLICATIONS & SYSTEMS TAPE ROADMAP.
- [4] 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*.
- [5] Sabaghian, K., Khamforoosh, K., & Ghaderzadeh, A. (2023). Data replication and placement strategies in distributed systems: A state of the art survey. *Wireless Personal Communications*, 129(4), 2419-2453.
- [6] 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*
- [7] Zheng, W., & Liu, W. (2025). Symmetry-Aware Transformers for Asymmetric Causal Discovery in Financial Time Series. *Symmetry*.
- [8] 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.
- [9] Tan, Y., Zheng, W., Cao, J., & Jiang, B. (2025). Intelligent Tax Systems: Automating Tax Audits and Improving Revenue Efficiency. *Open Journal of Accounting*, 14(3), 156-169.
- [10] Zhang, H., Ge, Y., Zhao, X., & Wang, J. (2025). Hierarchical Deep Reinforcement Learning for Multi-Objective Integrated Circuit Physical Layout Optimization with Congestion-Aware Reward Shaping. *IEEE Access*
- [11] Hu, X., Guo, L., Wang, J., & Liu, Y. (2025). Computational fluid dynamics and machine learning integration for evaluating solar thermal collector efficiency-Based parameter analysis. *Scientific Reports*, 15(1), 24528.

- [12] Elshazly, H., Ejarque, J., & Badia, R. M. (2022). Storage-heterogeneity aware task-based programming models to optimize I/O intensive applications. *IEEE Transactions on Parallel and Distributed Systems*, 33(12), 3589-3599.
- [13] Hoseinzadeh, M. (2019). A survey on tiering and caching in high-performance storage systems. *arXiv preprint arXiv:1904.11560*.
- [14] Cao, W., Mai, N., & Liu, W. (2025). Adaptive Knowledge Assessment via Symmetric Hierarchical Bayesian Neural Networks with Graph Symmetry-Aware Concept Dependencies. *Symmetry*.
- [15] Singh, G., Nadig, R., Park, J., Bera, R., Hajinazar, N., Novo, D., ... & Mutlu, O. (2022, June). Sibyl: Adaptive and extensible data placement in hybrid storage systems using online reinforcement learning. In *Proceedings of the 49th Annual International Symposium on Computer Architecture* (pp. 320-336).
- [16] Felten, F., Talbi, E. G., & Danoy, G. (2024). Multi-objective reinforcement learning based on decomposition: A taxonomy and framework. *Journal of Artificial Intelligence Research*, 79, 679-723.
- [17] Adebayo, Y. O., Adeusi, O. O., Adjadeh, J. P., Obion, S. M., & Abdulsalam, R. O. (2025). Artificial intelligence and predictive analytics to develop evidence-based migration policies for optimal integration and economic empowerment of migrants. *World Journal of Advanced Research and Reviews*, 25(1), 2147-2155.
- [18] Felsy, C., & Isaac Sajan, R. (2025). Actor-Critic Optimization Based Adaptive Task Scheduling Using Deep Double Dueling Q Network for Heterogeneous Cloud Environments. *Transactions on Emerging Telecommunications Technologies*, 36(7), e70185.
- [19] Zheng, W., Tan, Y., Jiang, B., & Wang, J. (2025). Integrating Machine Learning into Financial Forensics for Smarter Fraud Prevention. *Technology and Investment*, 16(3), 79-90.
- [20] 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.
- [21] Noon, N. (2024). Maximize Efficiency of Data Transfer Scheduling in Multi-Tiered Storage (Doctoral dissertation, University of Wollongong).
- [22] Pang, L., Alazzawe, A., Ray, M., Kant, K., & Swift, J. (2023). Adaptive Intelligent Tiering for modern storage systems. *Performance Evaluation*, 160, 102332.
- [23] Eytan, O., Harnik, D., Ofer, E., Friedman, R., & Kat, R. (2020). It's Time to Revisit {LRU} vs.{FIFO}. In *12th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage 20)*.
- [24] Kumar, S. S., & Agarwal, S. (2024). Rule based complex event processing for IoT applications: Review, classification and challenges. *Expert Systems*, 41(9), e13597.
- [25] Reyes, J., Solano-Charris, E., & Montoya-Torres, J. (2019). The storage location assignment problem: A literature review. *International Journal of Industrial Engineering Computations*, 10(2), 199-224.
- [26] Saxena, D., Kumar, J., Singh, A. K., & Schmid, S. (2023). Performance analysis of machine learning centered workload prediction models for cloud. *IEEE Transactions on Parallel and Distributed Systems*, 34(4), 1313-1330.
- [27] Drake, R. E., & Bond, G. R. (2023). Individual placement and support: History, current status, and future directions. *Psychiatry and Clinical Neurosciences Reports*, 2(3), e122.
- [28] 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.

- [29] Rasool, M. H., Taylan, O., Perwez, U., & Batunlu, C. (2023). Comparative assessment of multi-objective optimization of hybrid energy storage system considering grid balancing. *Renewable Energy*, 216, 119107.
- [30] Masdari, M., & Khoshnevis, A. (2020). A survey and classification of the workload forecasting methods in cloud computing. *Cluster Computing*, 23(4), 2399-2424.
- [31] Mienye, I. D., Swart, T. G., & Obaido, G. (2024). Recurrent neural networks: A comprehensive review of architectures, variants, and applications. *Information*, 15(9), 517.
- [32] Xing, S., & Wang, Y. (2025). Proactive Data Placement in Heterogeneous Storage Systems via Predictive Multi-Objective Reinforcement Learning. *IEEE Access*.
- [33] Sonkoly, B., Czentye, J., Szalay, M., Németh, B., & Toka, L. (2021). Survey on placement methods in the edge and beyond. *IEEE Communications Surveys & Tutorials*, 23(4), 2590-2629.