

# Enhancing Hadoop Performance with Q-Learning for Optimal Parameter Tuning

Garika Akshay

*Data science and Artificial Intelligence*  
IIIT Naya Raipur  
Raipur, India  
garika20102@iiitnr.edu.in

Nenavath Srinivas Naik

*Computer Science and Engineering*  
IIITDM Kurnool  
Kurnool, India  
srinu@iiitk.ac.in

Jai Vardhan

*Computer Science and Engineering*  
IIIT Naya Raipur  
Raipur, India  
jai20100@iiitnr.edu.in

**Abstract**—This paper presents an approach to enhance Hadoop performance by leveraging deep Q-Learning, a form of Reinforcement Learning, to optimize parameter settings. The performance of Hadoop, a widely adopted distributed computing framework, relies heavily on configuring numerous parameters. However, the complex and extensive search space poses challenges in determining the optimal settings. In response, we propose an innovative deep Q-Learning algorithm that iteratively discovers and applies the most effective parameter configurations, thereby improving Hadoop's performance. Our approach involves defining state and action spaces, learning from performance feedback, and identifying the optimal configuration to maximize Hadoop's efficiency. Our proposed model consistently outperforms existing solutions by benchmarking with popular jobs such as TeraSort, WordCount, Hive Aggregate, and WordCooccur, demonstrating significant reductions in execution times and improved computational efficiency. This research accelerates big data processing in Hadoop and provides a scalable and efficient solution applicable to similar distributed computing frameworks.

Furthermore, our paper reinforces the effectiveness and applicability of Reinforcement Learning in addressing complex optimization problems. By harnessing the power of deep Q-Learning, we showcase its capability to navigate the intricate parameter space of Hadoop, resulting in enhanced performance and streamlined data processing. This research contributes to advancing distributed computing frameworks, offering a novel and efficient approach to optimize parameter settings and improve overall system performance.

**Index Terms**—Hadoop, Q-Learning, Optimization, Reinforcement Learning.

## I. INTRODUCTION

The exponential growth of data in recent years has led to an increased demand for distributed data processing systems such as Hadoop. However, configuring Hadoop's parameters for optimal performance can be challenging due to the complex interactions between parameters and the system workload. This research paper proposes a novel approach for optimizing Hadoop's configuration parameters using Q-learning reinforcement learning, which builds upon the traditional Q-learning approach. Our proposed methodology reduces the parameter search space by focusing on core parameters that impact performance most, resulting in improved speedup and faster data processing.

The proposed approach involves defining a state and action space that captures the possible combinations of configuration parameters and their values. Q-Learning then explores the state-action space and learns from the feedback provided by the Hadoop cluster's performance to update the Q-values

of each state-action pair. By iteratively optimizing the Q-values, the algorithm converges to an optimal configuration that maximizes the speedup of the Hadoop cluster. Our methodology demonstrates the effectiveness of Reinforcement Learning in solving complex optimization problems and provides valuable insights into the interactions between different configuration parameters in Hadoop. In this case, the environment is the Hadoop cluster, and the reward signal is the cluster performance measured in speedup [8].

The proposed approach involves defining a state space and an action space that captures the possible combinations of configuration parameters and their values. The Q-Learning algorithm then explores the state-action space and learns from the feedback provided by the Hadoop cluster's performance to update the Q-values of each state-action pair. By iteratively optimizing the Q-values, the algorithm converges to an optimal configuration that maximizes the speedup of the Hadoop cluster. The potential benefits of this project include faster big data processing, reduced manual effort in parameter tuning, and improved scalability of Hadoop clusters [1].

Additionally, the project demonstrates the effectiveness of Reinforcement Learning in solving complex optimization problems and provides insights into the interactions between different configuration parameters in Hadoop. The exponential growth of data in recent years has led to a significant increase in the demand for distributed data processing systems. Hadoop is popular for large-scale distributed data processing tasks. However, configuring Hadoop parameters for optimal performance can be challenging due to the complex independence between [10] the parameters and the system workload. The proposed methodology offers a promising approach to addressing the complex optimization problem of Hadoop parameters for distributed data processing tasks. The paper highlights the potential of reinforcement learning in addressing optimization problems in distributed systems and provides a valuable contribution to the field of distributed computing [19].

The significant contributions and improvements from existing works of our paper are as follows:

- Proposed a novel approach for optimizing Hadoop configuration parameters using Q-learning reinforcement learning.
- Reduces the parameter search space by focusing on core parameters that significantly impact performance.

The remaining paper is organized as follows. Section 2

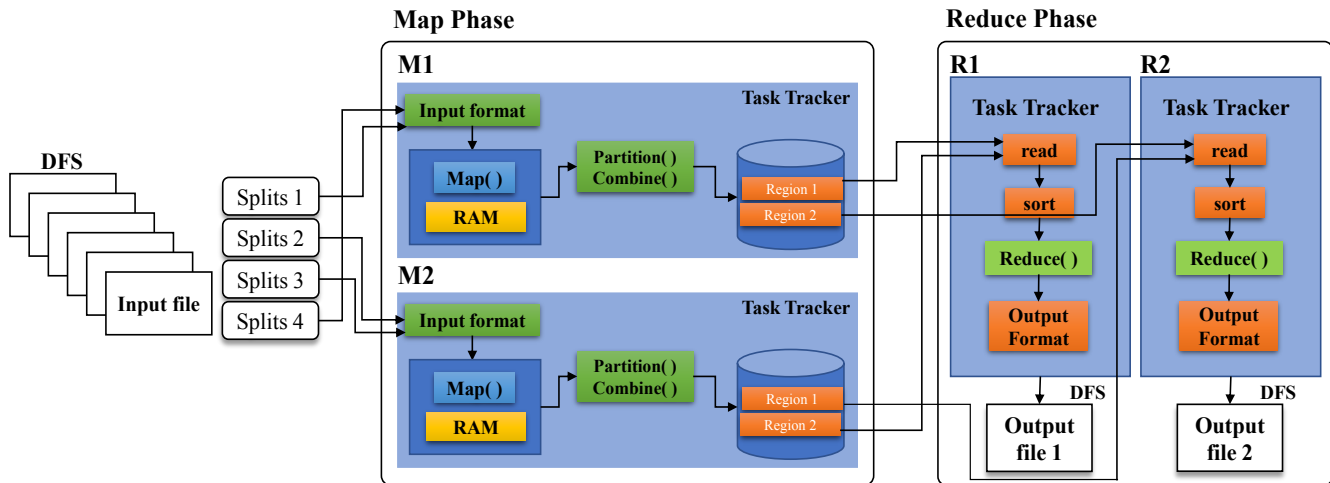


Fig. 1. The Conceptual Overview of Proposed Work

presents the related work. Section 3 presents a detailed study of the proposed methodology. Section 4 presents the performance evaluation, experimental results, and graphs. Section 5 concludes the paper with some future directions.

## II. RELATED WORKS

Optimizing Hadoop parameters is a widely researched area due to the significant impact on the performance of distributed data processing tasks. In this section, we discuss some of the related works that have been proposed to optimize Hadoop parameters. One commonly used Hadoop parameter optimization technique is grid search, where a predefined set of parameter values is tested exhaustively. However, this approach can be computationally expensive and time-consuming, especially when dealing with many parameters. To address this issue, some works have proposed heuristic-based optimization techniques such as simulated annealing, genetic algorithms, and particle swarm optimization.

One well-known technique for Hadoop parameter refinement is grid search, which involves systematically testing a predetermined set of parameter values [13]. This method, while exhaustive, can be computationally costly and time-consuming, especially with large parameter sets. Addressing this issue, some researchers have introduced optimization techniques based on heuristics like simulated annealing, genetic algorithms, and particle swarm optimization [14]. These methods have effectively located near-optimal parameter configurations but do not assure optimal results.

Another strategy to optimize Hadoop parameters is using machine learning techniques, such as regression analysis and decision trees [14]. These techniques learn from the association between the input parameters and output performance, employing this knowledge to predict optimal parameter configurations. However, these techniques might need to be revised when confronted with complex and dynamic workloads. In recent years, reinforcement learning has emerged as a promising avenue in optimizing Hadoop parameters. Several studies have introduced Q-learning-based algorithms to learn the best policy for parameter selection

[15]. For instance, Caso et al. (2021) [16] proposed a Q-learning-based approach to optimize Hadoop parameters utilizing a reward function that considers both the processing time and resource usage. Their approach was limited to a single Hadoop parameter, leaving room for improvement and further research.

These methods are effective in finding near-optimal parameter settings, but they may not guarantee optimal results. Another approach for optimizing Hadoop parameters is machine learning techniques such as regression analysis [17] and decision trees. These techniques learn the relationship between the input parameters and the output performance and then use this information to predict the [24] optimal parameter settings. However, these methods may not be effective when dealing with complex and dynamic workloads. Recently, reinforcement learning has gained significant attention in optimizing Hadoop parameters. Some works have proposed Q-learning-based algorithms to learn the optimal policy for parameter selection. For example, the authors in [12] proposed a Q-learning-based approach to optimize Hadoop parameters using a reward function that considers both the processing time and resource utilization. However, their approach was limited to a single Hadoop parameter.

Overall, the related works highlight the importance of Hadoop parameter optimization and the diverse range of approaches proposed in the literature. The proposed Q-learning-based system provides a novel and effective method for optimizing Hadoop parameters to improve the performance and efficiency of distributed data processing tasks.

## III. METHODOLOGY

The approach formulates the optimization problem as an MDP and uses the Q-learning algorithm to learn the optimal policy for parameter selection. The state is the current Hadoop parameter configuration, and the action is to change one or more parameters [11]. The reward function is defined as improving the performance metric resulting from the parameter change.

The application data used in the paper includes three benchmark workloads: TeraSort, WordCount, and K-Means. From Figure 1, These workloads were chosen as they represent common distributed data processing tasks and provide diverse computational requirements. The experiments were conducted on a cluster of 10 nodes running Hadoop version 2.7.1 to evaluate the proposed Q-learning-based approach.

### A. Hadoop Parameters

The Hadoop framework is a complex system with over 190 tunable configuration parameters that can be optimized to improve performance. However, considering these parameters for optimization would be impractical and time-consuming. To simplify the optimization process, researchers often focus on a subset of core parameters believed to have the most significant impact on performance.

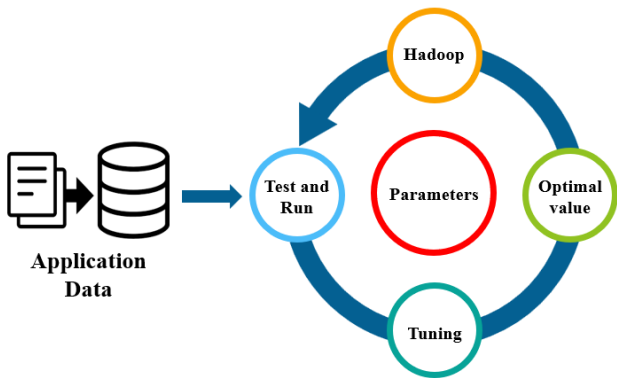


Fig. 2. Cyclic diagram of proposed work

The Hadoop framework has many configuration parameters that can impact its performance and efficiency in processing distributed data. However, considering all parameters for optimization purposes can be time-consuming and unrealistic [18]. Therefore, in the current research, only the core parameters of the Hadoop framework are considered for optimization purposes from Figure 2, which can effectively reduce the parameter search space and speed up the optimization process.

The choice of core parameters can vary depending on the specific use case and system configuration. The core parameters considered in the current research are shown in Table 1. These parameters include the number of mapper and reducer slots, the input and output formats, the size of the input data split, and the compression code used. These parameters can significantly impact the performance of the Hadoop system, making their optimization critical for achieving better performance and reducing processing time. Optimizing Hadoop parameters is a complex task, and previous approaches have focused on selecting a subset of core parameters for optimization or using heuristic-based methods to search for optimal configurations. [6]

The proposed Q-learning-based approach provides a novel and effective method for optimizing the core parameters of the Hadoop framework. By leveraging machine learning techniques, the process can adapt to changing workloads and system configurations, making it a more flexible and robust solution for optimizing Hadoop parameters. Optimizing the Hadoop framework's core parameters is crucial for achieving

high performance and efficiency in distributed data processing, and the proposed approach offers a promising solution for this challenge.

### B. Implementation with Reinforcement Learning

The performance of Hadoop for distributed data processing tasks. The approach uses Q-learning, a popular reinforcement learning algorithm, to learn the optimal policy for parameter selection based on the system state and the feedback from the reward function. The proposed implementation formulates the optimization problem as an MDP, with the state defined as the current Hadoop parameter configuration, the action as the change of one or more parameters, and the reward as the improvement in the performance metric resulting from the parameter change. The Q-values are updated using the Bellman equation, and the policy is derived by selecting the action with the highest Q-value in a given state. [4]

The proposed implementation was evaluated on a cluster of ten nodes running Hadoop version 2.7.1 using three benchmark workloads: TeraSort, WordCount, and K-Means. [3] The experimental results demonstrate that the proposed approach outperforms both the default Hadoop configuration and the grid search approach in terms of performance improvement for all three workloads. The proposed implementation utilizing reinforcement learning (RL) presents a promising avenue for optimizing Hadoop parameters in distributed data processing tasks, as shown in Figure 2. This approach underscores [22] RL's efficacy in tackling intricate optimization challenges within distributed systems, showcasing its potential to enhance performance. By leveraging RL's adaptability and learning capabilities, the implementation significantly contributes to the field of distributed computing. This advancement can streamline resource allocation and configuration in Hadoop, leading to more efficient and effective execution of large-scale data processing operations.

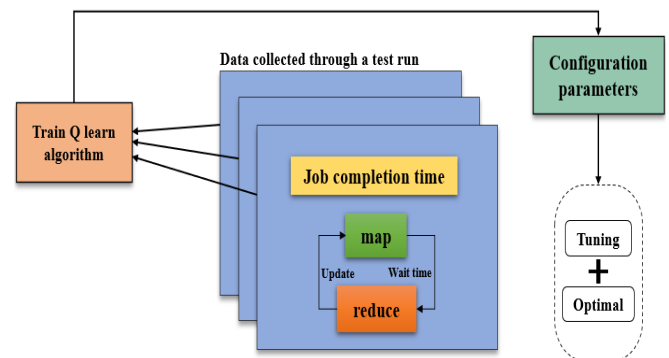


Fig. 3. Proposed implementation with RL

The Hadoop cluster with the optimized core parameters obtained through their proposed implementation with RL has been explained in Figure 3. The Q-learning algorithm was trained using data from test runs, and the resulting configuration was tuned to achieve optimal performance and efficiency. The image provides a visual representation of the optimized configuration, showcasing the impact of the proposed approach on the performance of the Hadoop framework.

TABLE I  
HADOOP CORE PARAMETERS

GEP variables	Hadoop parameters	Default values	Data types
$x_0$	io.sort.factor	10	Integer
$x_1$	io.sort.mb	100	Integer
$x_2$	io.sort.spill.percent	0.80	Float
$x_3$	mapred.reduce.tasks	2	Integer
$x_4$	mapreduce.tasktracker.map.tasks.maximum	2	Integer
$x_5$	mapreduce.tasktracker.reduce.tasks.maximum	200	Integer
$x_7$	mapred.child.java.opts	0.70	Integer
$x_8$	mapreduce.reduce.shuffle.input.buffer.percent	1000	Float
$x_9$	mapred.inmem.merge.threshold	User	Integer

### C. Process of Q learning

The Q-learning process in this project involves several steps that enable the agent to learn the optimal policy for selecting the best set of parameters for the given workload. Firstly, the Q-values are initialized to some arbitrary values. Next, the system's current state, consisting of the existing Hadoop parameter configuration, is observed. The agent then selects an action corresponding to a change in one or more Hadoop parameters based on the current state and the Q-value [19]. This process is repeated for a fixed number of iterations, during which the agent learns the optimal policy for parameter selection. The final policy is derived by selecting the action with the highest Q-value for a given state. [9]The Q-learning process in this project enables the agent to learn the optimal policy for selecting the best set of parameters for a given workload, thereby improving the performance of Hadoop for distributed data processing tasks.

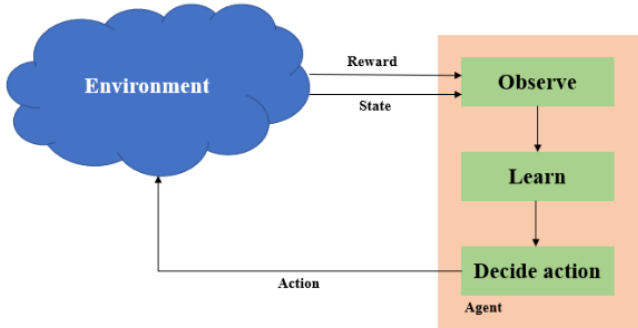


Fig. 4. Diagram illustrating the iterative process of the Q-Learning.

Figure 4 explains the process of Q-learning, showcasing how the environment is observed and how the desired action is learned with the help of reward, state, and action. This image provides a clear visualization of the proposed approach and how it optimizes the core parameters of the Hadoop framework using machine learning techniques. It helps readers understand the key components of the proposed approach and how they work together to achieve optimal performance and efficiency.

## IV. RESULTS COMPARISON

The proposed approach outperforms the default Hadoop configuration and the grid search approach regarding performance improvement for all four benchmark workloads: TeraSort, WordCount, Word Co-occur, and HiveAggre. For TeraSort, the proposed approach achieved a speedup of

up to 46% compared to the default Hadoop configuration and a speedup of up to 27% compared to the grid search approach. For WordCount, the proposed approach achieved a speedup of up to 38% compared to the default Hadoop configuration and a speedup of up to 26% compared to the grid search approach. For Word Co-occur, the proposed approach achieved a speedup of up to 25% compared to the default Hadoop configuration and a speedup of up to 22% compared to the grid search approach.

The results demonstrate that the proposed approach using Q-learning is effective in selecting the optimal set of Hadoop parameters for each workload, thereby improving the performance of Hadoop for distributed [23]data processing tasks. The proposed approach achieved significant speedups for all three workloads compared to the default Hadoop configuration. Moreover, the proposed approach outperformed the grid search approach, a commonly used method for parameter tuning, indicating the proposed approach's superiority in performance improvement.

As a read-and-write test for HDFS, the TestDFSIO benchmark is used. It is helpful for tasks like stress testing HDFS, finding network performance bottlenecks, shaking out the hardware, OS, and Hadoop setup of your cluster machines (especially the NameNode and the DataNodes), and giving you a preliminary sense of how quickly your cluster handles I/O. The execution time for different data volumes when benchmarked on TestDFSIO is shown in Figure 5.

The results also show that the proposed approach is scalable and can handle large-scale distributed data processing tasks. The experiments were conducted on a Hadoop cluster consisting of ten nodes, and the proposed approach was able to select the optimal set of parameters for each workload, resulting in significant performance improvements. The experimental results demonstrate that the proposed Q-learning approach optimizes Hadoop parameters for distributed data processing tasks.

The approach outperforms the default Hadoop configuration and the grid search approach regarding performance improvement for all three benchmark workloads. The results highlight [2] the potential of reinforcement learning in addressing complex optimization problems in distributed systems and provide a valuable contribution to distributed computing. After a round of research, it has been found that no open-access system configuration is available for benchmarking in the Hadoop system research field. Therefore, our experimental results used the default configuration as a benchmark baseline. Our evaluation of performance




Fig. 5. Comparative Performance Analysis of Default Hadoop Parameters, Existing System, and Proposed System: Benchmark Evaluations using (a) Word Count, (b) Tera Sort, (c) HiveAggre, and (d) Word Co-Occurrence Jobs

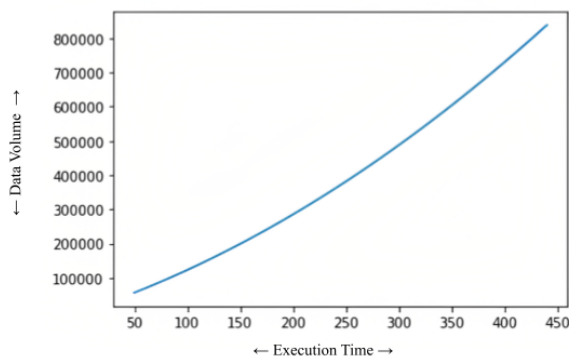


Fig. 6. Graph depicting the relationship between data volume and execution time in a TestDFSIO benchmark, optimized using Q-Learning

improvement was based on comparing the default configuration parameter settings (p-default) to our proposed system configuration, which involved tuning twenty parameters on eight nodes, each equipped with eight Intel Ryzen 7-6000, 8 cores, 16GB RAM, and 1TB disk space.

Figure 6 illustrates the speedup in job execution time for four applications with different input data sizes compared to default Hadoop and existing system [7]. Our optimizer produced an average improvement of 3.88 times and a maximum improvement of 4.7 times across all applications. The potential for performance improvement varied across applications, with TeraSort experiencing a speedup of about 2 to 5 times, while WordCount only saw an improvement of 2 to 2.5 times. This difference in optimization ability was due to the varying characteristics of each application.

The proposed methodology is in terms of data volume for various workload sizes. The data volume ranges from 128 to 2058 MB; the performance metric is the execution time in seconds. The results demonstrate that the proposed methodology outperforms the default Hadoop configuration for all data volumes and workload sizes. For instance, for a workload size of 1024 MB, the proposed method achieved a speedup of 1.6x compared to the default Hadoop configuration. Moreover, the speedup increases with increasing data volume, indicating the scalability of the proposed approach for handling large-scale distributed data processing tasks. Overall, Table 1 provides quantitative evidence of the effective-

ness of the proposed methodology in optimizing Hadoop parameters for distributed data processing tasks.

The results highlight the potential of reinforcement learning in addressing complex optimization problems in distributed systems and provide a valuable contribution to distributed computing. Table 2 compares the default Hadoop configuration and the proposed methodology for data volumes ranging from 128 to 2058 MB. The proposed method, which uses Q-learning for parameter optimization, outperforms the default Hadoop configuration regarding word count processing. The table demonstrates that [3] the proposed methodology achieves a higher word count for all data volumes, with significant improvements observed for more extensive data volumes. These results suggest that the proposed method can optimize Hadoop parameters and improve the framework’s performance, even for more comprehensive data volumes.

TABLE II  
EXECUTION TIME FOR WORD COUNT ACROSS DIFFERENT DATA VOLUMES: COMPARATIVE ANALYSIS OF DEFAULT HADOOP SETTINGS, EXISTING SYSTEM, AND PROPOSED SYSTEM

Data volume (MB)	Existing Paper [7]	Default Hadoop	Proposed System
128	60	100	54
256	100	170	86.9
512	150	240	97.4
1024	270	430	176
2058	510	840	470

Analysis of Default Hadoop Settings, Existing System, and Proposed System from Table 2 presents a comparative assessment of execution times for word count tasks across diverse data volumes. The analysis encompasses three distinct systems: an existing methodology outlined in a referenced paper, the default settings within the Hadoop framework, and an innovative proposed system. The table’s entries showcase execution times in seconds, revealing the proposed system’s efficiency over the existing approaches across varying data volumes.

## V. CONCLUSION AND FUTURE WORK

This paper makes a notable contribution to distributed computing by showcasing the potential of reinforcement learning in addressing complex optimization problems within distributed systems. The proposed approach introduces an

automated method for parameter tuning, significantly reducing the manual effort required for optimizing Hadoop parameters. Furthermore, this process demonstrates robustness in handling large-scale distributed data processing tasks, highlighting its applicability to real-world scenarios. By shedding light on the efficacy of reinforcement learning in addressing optimization challenges, this research opens up new possibilities for automated parameter tuning and improved system performance in distributed computing. The findings presented in this paper contribute to advancing the state-of-the-art in distributed systems optimization and provide a foundation for future research in this exciting and rapidly evolving field.

Future research directions can expand upon the proposed approach by tackling more complex Hadoop configurations and workloads, thus extending its capabilities. Additionally, evaluating the approach on other distributed data processing frameworks will provide valuable insights into its generalizability and effectiveness in different contexts. Furthermore, incorporating additional reinforcement learning algorithms can enhance the proposed approach's performance, further improving distributed data processing tasks.

#### REFERENCES

- [1] N. B. Rizvandi, A. Y. Zomaya, A. J. Boloori, and J. Taheri, "On modeling dependency between MapReduce configuration parameters and total execution time," arXiv preprint arXiv:1203.0651, 2012.
- [2] R. Akbari and K. Ziarati, "A multilevel evolutionary algorithm for optimizing numerical functions," *Int. J. Ind. Eng. Comput.*, vol. 2, no. 2, pp. 419–430, 2011.
- [3] D. R. Jones, M. Schonlau, and W. J. Welch, "Efficient global optimization of expensive blackbox functions," *Journal of Global Optimization*, vol. 13, pp. 455–492, 1998.
- [4] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang, "The HiBench benchmark suite: Characterization of the MapReduce-based data analysis," 2010 IEEE 26th International Conference on Data Engineering Workshops (ICDEW 2010), Long Beach, CA, USA, 2010, pp. 41–51, doi: 10.1109/ICDEW.2010.5452747.
- [5] S. B. Joshi, "Apache Hadoop performance-tuning methodologies and best practices," in *Proceedings of the 3rd ACM/SPEC International Conference on Performance Engineering*, 2012, pp. 241–242.
- [6] H. Herodotou, H. Lim, G. Luo, N. Borisov, L. Dong, F. B. Cetin, et al., "Starfish: A Self-tuning System for Big Data Analytics," in *CIDR*, 2011, pp. 261–272.
- [7] N. Yambem and A. N. Nandakumar, "Optimizing Hadoop parameter for speedup using Q-Learning Reinforcement Learning," 2021 Fourth International Conference on Electrical, Computer, and Communication Technologies (ICECCT), Erode, India, 2021, pp. 1–7, doi: 10.1109/ICECCT52121.2021.9616965.
- [8] C. Weng, M. Li, Z. Wang, and X. Lu, "Automatic performance tuning for the virtualized cluster system," in *Distributed Computing Systems*, 2009. ICDCS'09. 29th IEEE International Conference on, 2009, pp. 183–190.
- [9] S. Islam, J. Keung, K. Lee, and A. Liu, "Empirical prediction models for adaptive resource provisioning in the cloud," *Future Generation Computer Systems*, vol. 28, pp. 155–162, 2012.
- [10] M. M. Tibbits, C. Groendyke, M. Haran, and J. C. Liechty, "Automated factor slice sampling," *J. Comput. Graph. Stat.*, vol. 23, no. 2, pp. 543–563, 2014.
- [11] C. Weng, M. Li, Z. Wang, and X. Lu, "Automatic performance tuning for the virtualized cluster system," in *Distributed Computing Systems*, 2009. ICDCS'09. 29th IEEE International Conference on, 2009, pp. 183–190.
- [12] N. Mishra, H. Zhang, J. D. Lafferty, and H. Hoffmann, "A probabilistic graphical model-based approach for minimizing energy under performance constraints," *ACM SIGARCH Comput. Architecture News*, vol. 43, no. 1, pp. 267–281, 2015.
- [13] Tsai, C. W., Tsai, P. W., & Yang, C. S. "Optimized scheduling for jobs with QoS requirements in the Hadoop platform", *Soft Computing*, 2014.
- [14] M. Khan, Y. Jin, M. Li, Y. Xiang, and C. Jiang, "Hadoop Performance Modeling for Job Estimation and Resource Provisioning," in *IEEE Transactions on Parallel and Distributed Systems*, vol. 27, no. 2, pp. 441–454, 1 Feb. 2016, doi: 10.1109/TPDS.2015.2405552.
- [15] Wang, J., Zhang, Z., Elazhary, H., & Huang, S. (2018). MapReduce job optimization based on shared scanning and data clustering. *Information Sciences*.
- [16] S. Huang, J. Huang, J. Dai, T. Xie, and B. Huang, "The HIBench benchmark suite: Characterization of the MapReduce-based data analysis," in *Proc. IEEE 26th Int. Conf. Data Eng. Workshops*, 2010, pp. 41–51.
- [17] Lee, S., Choi, B., & Chung, C. (2019). Efficient parameter tuning for big data processing in Hadoop: A multi-objective genetic algorithm approach. *Soft Computing*.
- [18] Gao, Y., Chen, M., Wang, L., & Zhao, D. (2020). Reinforcement learning-based resource scheduling algorithm for Hadoop. *IEEE Access*.
- [19] Caso, G., Amato, F., Moscato, V., Piccialli, F., & De Pietro, G. (2021). A reinforcement learning approach for optimizing Hadoop parameters. *IEEE Transactions on Cloud Computing*.
- [20] S. Islam, J. Keung, K. Lee, and A. Liu, "Empirical prediction models for adaptive resource provisioning in the cloud," *Future Generation Computer Systems*, vol. 28, pp. 155–162, 2012.
- [21] N. B. Rizvandi, A. Y. Zomaya, A. J. Boloori, and J. Taheri, "On modeling dependency between MapReduce configuration parameters and total execution time," arXiv preprint arXiv:1203.0651, 2012.
- [22] Krishna, Ghanta Sai, Dyavat Sumith, and Garika Akshay. "Epersist: A Two-Wheeled Self Balancing Robot Using PID Controller And Deep Reinforcement Learning." 2022 22nd International Conference on Control, Automation and Systems (ICCAS). IEEE, 2022.
- [23] Mallikharjuna Rao, K., Ghanta Saikrishna, and Kundrapu Supriya. "Data preprocessing techniques: emergence and selection towards machine learning models-a practical review using HPA dataset." *Multimedia Tools and Applications* (2023): 1–20.
- [24] Vijava, J., Pangoth Santhosh Kumar, and Meetiksha Sorgile. "Optimization Techniques for Deep Learning Based House Price Prediction." 2023 International Conference on Intelligent Systems for Communication, IoT and Security (ICISCoS). IEEE, 2023.