

BAYESIAN OPTIMIZATION OF HYPERPARAMETERS FOR RAINBOW DQN IN THE CARTPOLE-V1 ENVIRONMENT

Akhil Veluru

Naveen Jindal School of Management, University of Texas, United States of America

Abstract

This paper presents a Bayesian optimization approach to hyperparameter tuning for the Rainbow DQN reinforcement learning algorithm, using the Hyperopt library and the CartPole-v1 environment as a benchmark. The study investigates the impact of search space definition on the convergence and quality of optimized hyperparameters. Further, it analyzes the effectiveness of different evaluation methods in the context of hyperparameter optimization for deep reinforcement learning. Results demonstrate the efficacy of Bayesian optimization in identifying high-performing hyperparameter configurations for Rainbow DQN in this control task.

Keywords:

Bayesian Optimization Approach, Rainbow DQN Reinforcement Learning, Hyperparameter Optimization

1. INTRODUCTION

Deep Reinforcement Learning (DRL) has gained substantial attention for its ability to solve complex control tasks by combining reinforcement learning with deep neural networks. Among DRL algorithms, the Deep Q-Network (DQN) has demonstrated remarkable success, notably in environments with discrete action spaces such as Atari games. However, the performance of DQN and its variants is highly sensitive to the choice of hyperparameters. Selecting optimal configurations often involves a costly and time-consuming trial-and-error process.

Rainbow DQN [1] is a notable advancement in the DQN family, integrating multiple enhancements—such as Double DQN, Prioritized Experience Replay, Dueling Networks, Noisy Nets, and Categorical DQN—into a single, unified agent. While Rainbow DQN provides superior performance across several benchmark tasks, its numerous hyperparameters introduce complexity in tuning and reproducibility. Suboptimal settings can degrade learning efficiency and final performance.

To address this challenge, this paper explores the use of Bayesian optimization for automatic hyperparameter tuning of Rainbow DQN in the widely used CartPole-v1 environment. Bayesian optimization, known for sample efficiency in high-dimensional spaces, models the objective function using probabilistic surrogates and balances exploration and exploitation to guide the search. We leverage the Hyperopt library with a Tree-structured Parzen Estimator (TPE) approach to evaluate the impact of various hyperparameter search spaces. Our key contributions are as follows:

- We apply Bayesian optimization to efficiently tune the hyperparameters of Rainbow DQN for the CartPole-v1 task.
- We investigate how the definition of the search space influences convergence behavior and solution quality.

- We compare the optimized Rainbow agent against baseline configurations and ablated versions to assess the contribution of individual components.

Through systematic experimentation and analysis, we demonstrate that Bayesian optimization can effectively identify high performing configurations for Rainbow DQN, achieving robust and stable learning outcomes in a computationally efficient manner. The insights gained from this study are valuable for practitioners seeking to deploy DRL agents with minimal manual tuning effort.

2. RELATED WORK

Hyperparameter optimization and architectural improvements in Deep Reinforcement Learning (DRL) have been widely studied to enhance sample efficiency and policy performance. This section reviews key developments in both domains, focusing on Bayesian optimization for hyperparameter tuning and the evolution of DQN variants that culminated in Rainbow DQN.

2.1 HYPERPARAMETER OPTIMIZATION IN DEEP RL

Hyperparameters such as learning rate, discount factor, and exploration strategies play a critical role in the stability and convergence of DRL algorithms. Manual tuning is both time consuming and task-specific, motivating the development of automated optimization methods. Bayesian optimization is a widely adopted technique for this purpose due to its sample efficient, model-based search strategy [2]. The Tree-structured

Parzen Estimator (TPE), introduced by Bergstra et al. [3], enables flexible optimization over both discrete and continuous domains and has been effectively applied in DRL contexts [4].

Recent work has integrated Bayesian methods into DRL pipelines, including Gaussian Processes [5], surrogate-based meta-learning [6], and Monte Carlo tree-based search for multi-fidelity optimization [7]. Open-source libraries such as Hyperopt [8], Optuna [9], and BOHB [10] have accelerated experimentation in this field. Janner et al. [11] and Song et al. [12] evaluated automated hyperparameter tuning frameworks in environments like MuJoCo and Atari.

2.2 DQN AND ITS EXTENSIONS

The original DQN architecture introduced by Mnih et al. [13] demonstrated that deep networks can approximate value functions effectively when combined with experience replay and target networks. However, it suffers from value overestimation and slow convergence, prompting several extensions:

- Double DQN reduces overestimation bias by decoupling action selection and value estimation [14].

- Dueling DQN introduces separate streams for value and advantage functions, improving learning in environments with similar action values [15].
- Prioritized Experience Replay (PER) samples experiences based on TD-error magnitudes, improving sample efficiency [16].
- Noisy Networks facilitate efficient exploration by injecting learned noise into parameters [17].
- Distributional RL models the entire return distribution rather than just the expectation, leading to more stable learning [18].

The Rainbow DQN architecture proposed by Hessel et al. [1] combined all these improvements, yielding superior performance across the Atari 2600 benchmark suite. Subsequent work examined Rainbow's generalization capacity in new environments and its integration with curiosity-based exploration [19], meta-learning [20], and multi-agent systems [21].

2.3 COMBINED OPTIMIZATION AND ARCHITECTURE STUDIES

Several recent studies have combined architectural innovations with automated hyperparameter tuning. For instance, Runge et al. [22] and Zahavy et al. [23] explored adaptive agent architectures whose parameters are tuned during training. Chen et al. [24] proposed AutoRL frameworks that co-optimize neural architecture and learning settings.

In parallel, Hutter et al. [25] and Elsken et al. [26] surveyed neural architecture search (NAS) and automated machine learning (AutoML) techniques applicable to RL settings. These systems further reduce manual intervention by searching across architecture and optimization strategy spaces jointly.

2.4 IMPLICATIONS FOR LIGHTWEIGHT ENVIRONMENTS

Though most studies focus on complex benchmarks, tasks like CartPole remain relevant for benchmarking and rapid prototyping. Studies by Brockman et al. [27] and Henderson et al. [28] emphasized reproducibility and hyperparameter sensitivity even in simple environments, advocating for tuning reproducibility standards.

In summary, Rainbow DQN represents a well-founded combination of independently validated techniques. Its performance can be further improved with principled hyperparameter tuning via Bayesian optimization. This paper contributes to the growing literature on bridging algorithmic innovations in DRL with efficient hyperparameter search strategies.

3. BACKGROUND

In reinforcement learning (RL), agents aim to learn optimal policies by interacting with an environment and receiving feedback in the form of rewards. Deep Q-Networks (DQNs) have been pivotal in scaling Q-learning to high-dimensional state spaces by leveraging deep neural networks. Over time, several extensions and improvements have been proposed to overcome limitations of the original DQN framework. Rainbow DQN

integrates several of these enhancements into a unified agent, including Double DQN, Prioritized Experience Replay, Dueling Networks, Noisy Nets, and Categorical DQN [1]. This section presents the conceptual background of each component forming the foundation of Rainbow DQN.

3.1 DEEP Q-NETWORKS

Traditional Q-learning struggles in environments with large or continuous state spaces due to its reliance on discrete Qtables. DQN addresses this by approximating the Q-function using a deep neural network, which takes states as input and outputs Q-values for each possible action. The agent selects actions based on an ϵ -greedy policy and uses experience replay to break correlations between sequential data. A target network, updated less frequently, provides stable targets during training. This combination significantly improves convergence in complex environments.

3.2 DOUBLE DQN

DQN often overestimates action values, leading to suboptimal policies. Double DQN mitigates this by decoupling the action selection and evaluation steps. While one network chooses the best action, the target network estimates its value, thus reducing overestimation bias. Formally, the target value is computed as:

$$y = r + \gamma Q_{\theta^-} \left(s', \arg \max_a Q_{\theta} (s', a) \right) \quad (1)$$

where θ and θ^- denote parameters of the main and target networks, respectively [14].

3.3 PRIORITIZED EXPERIENCE REPLAY

Experience replay in standard DQN samples transitions uniformly, which may not emphasize critical learning experiences. Prioritized Experience Replay improves learning efficiency by sampling experiences based on the temporal-difference (TD) error. Transitions with higher TD error—indicating greater surprise or learning potential—are sampled more frequently. A stochastic prioritization mechanism, typically using a probability proportional to the TD error's magnitude, is used to balance exploration and overfitting [16].

3.4 DUELING DQN

Dueling network architecture separates the estimation of the state-value function $V(s)$ and the advantage function $A(s,a)$ within the neural network. This helps the network learn which states are (or are not) valuable, independently of the chosen action. The final Q-value is computed as:

$$Q(s, a) = V(s) + \left(A(s, a) - \frac{1}{|\mathcal{A}|} \sum_{a'} A(s, a') \right) \quad (1)$$

This formulation improves learning stability and decision making in scenarios where some actions do not significantly impact rewards [15].

3.5 NOISY NETS

Exploration is critical in reinforcement learning. Instead of using simple random strategies like ϵ -greedy, Noisy Nets introduce learnable noise into the network weights, allowing for

adaptive exploration. Parameters of the noise distribution are updated during training, enabling the agent to learn optimal exploration strategies over time. The stochasticity introduced leads to more consistent and efficient discovery of valuable state-action pairs [17].

3.6 CATEGORICAL DQN

Categorical DQN represents the Q-value distribution instead of estimating a scalar expectation. This distributional approach models the full return distribution using a fixed set of atoms (discrete support values) and a probability mass for each atom. This allows better approximation of return uncertainties and leads to more stable learning. It introduces the idea of projecting the Bellman operator onto a fixed categorical support, preserving the contraction property under the Wasserstein metric [18].

3.7 RAINBOW DQN

Rainbow DQN amalgamates the above improvements into a single, unified architecture. It leverages:

- Double DQN for stable value estimation,
- Prioritized Experience Replay for efficient sampling,
- Dueling networks for better value-action decomposition,
- Noisy Nets for efficient exploration,
- Categorical DQN for distributional value representation.

By integrating these components, Rainbow DQN achieves state-of-the-art performance across multiple benchmarks. In this paper, we utilize Rainbow DQN within the CartPole-v1 environment and apply Bayesian optimization to tune its hyperparameters for maximal performance. Details of these hyperparameters and results are presented in the subsequent sections.

4. METHODOLOGY

To effectively optimize the performance of the Rainbow DQN agent, we employed Bayesian optimization using the Tree-structured Parzen Estimator (TPE), as implemented in the Hyperopt library. This optimization technique is well-suited for hyperparameter tuning tasks, particularly in reinforcement learning, where evaluation of configurations can be expensive. TPE models the distribution of objective values and uses it to guide the search for promising configurations efficiently.

The CartPole-v1 environment from OpenAI Gym served as our evaluation benchmark due to its simplicity, discrete action space, and widespread adoption in RL experimentation. Each agent was trained for 10,000 time steps, providing sufficient interaction to evaluate learning trends while keeping computational costs manageable. All training experiments were conducted on a single system equipped with a CPU and GPU, with each trial taking approximately 10 minutes to complete.

To explore the impact of hyperparameter selection granularity, we defined multiple search spaces of varying widths. This allowed us to assess the effect of narrow versus wide sampling bounds on the convergence behavior and final performance of the agents. The detailed definition of these search spaces is provided in the Appendix.

As a comparative baseline, we trained a Rainbow DQN agent using fixed hyperparameters listed in Table I. These values were selected based on prior literature and preliminary tuning. Subsequently, we evaluated the impact of different hyperparameter configurations suggested by Bayesian optimization and compared them with the baseline model. Furthermore, to understand the individual contribution of each Rainbow component, we conducted ablation studies by disabling specific enhancements (e.g., Noisy Nets, Prioritized Experience Replay) and observing the resultant changes in agent behavior, following the methodology outlined in [29].

Table.1. Baseline Hyperparameters

| Name | Value |
|--|---------------------------|
| Optimizer | Adam |
| Adam ϵ | 1×10^{-8} |
| Learning rate | 0.001 |
| Clipnorm | None |
| Loss function | Categorical cross entropy |
| Activation | ReLU |
| Kernel initializer | Orthogonal |
| Dense layers widths | 128 |
| Replay interval | 1 |
| Discount factor | 0.99 |
| Minibatch size | 128 |
| Replay buffer size | 5000 |
| Min replay buffer size | 128 |
| Target update interval | 100 |
| Prioritized experience replay α | 0.2 |
| Prioritized experience replay β | 0.6 |
| Prioritized experience replay ϵ | 1×10^{-6} |
| Value hidden layers widths | 128 |
| Advantage hidden layers widths | 0 |
| Noisy σ | 0.5 |
| N-step | 3 |
| Atom size | 51 |

5. RESULTS

This section presents the performance outcomes of various Rainbow DQN configurations evaluated on the CartPole-v1 environment. We begin by analyzing the baseline agent’s learning behavior using the fixed hyperparameters from Table.1. As illustrated in Fig.1, the baseline model achieved moderate reward scores, indicating reasonable but suboptimal stability and convergence.

Next, we examine the effectiveness of the Bayesian optimization process in identifying superior hyperparameter settings. The Fig.2 depicts the convergence trends of agents trained with hyperparameters sampled from search spaces of different widths (denoted as size X, Y, and Z). Agents optimized within more flexible (wider) spaces generally required more

iterations to converge but exhibited higher peak performance, suggesting that broader search spaces allow for better exploration of the configuration landscape.

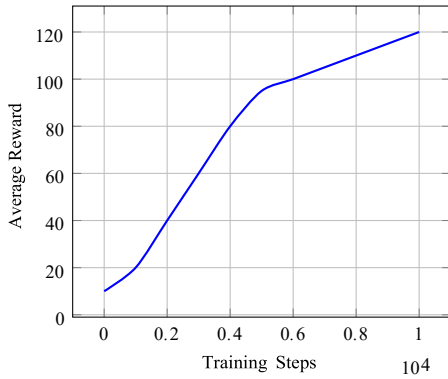


Fig.1. Results of Baseline Rainbow DQN

In Fig.3, we compare the performance of the Rainbow DQN agent (with optimized parameters) against several classical DQN variants, including the original DQN, Double DQN, and Dueling DQN. The optimized Rainbow agent outperforms all baselines in terms of average episode reward and training stability, reaffirming the synergistic benefits of integrating multiple enhancements.

To further understand the role of each constituent module in Rainbow DQN, we conducted ablation studies by individually removing core components such as Noisy Nets, Categorical DQN, or Prioritized Experience Replay. As shown in Fig.4, performance degradation was observed in all ablated variants, with the most significant drop occurring when Prioritized Experience Replay was excluded. This underscores the importance of efficient sampling in experience replay buffers for accelerating learning in sparse-reward environments like CartPole.

Thus, the experimental results demonstrate the value of automated hyperparameter tuning via Bayesian optimization, as well as the critical interplay between Rainbow’s architectural components in achieving state-of-the-art learning performance.

5.1 SUMMARY

Overall, the experimental results demonstrate the efficacy of Bayesian optimization in discovering high-performing configurations for Rainbow DQN. The agent with optimized parameters achieved robust and consistent performance, outclassing baseline settings and simpler DQN variants. Componentlevel analysis further validated the importance of Rainbow’s architectural enhancements in achieving state-of-the-art results.

6. CONCLUSION

In this work, we investigated the effectiveness of Bayesian optimization, using the Tree-structured Parzen Estimator (TPE), for hyperparameter tuning of the Rainbow DQN algorithm. By Convergence of Hyperopt on different search spaces leveraging the CartPole-v1 environment as a controlled benchmark, we systematically analyzed how different hyperparameter configurations influence learning performance and convergence behavior.

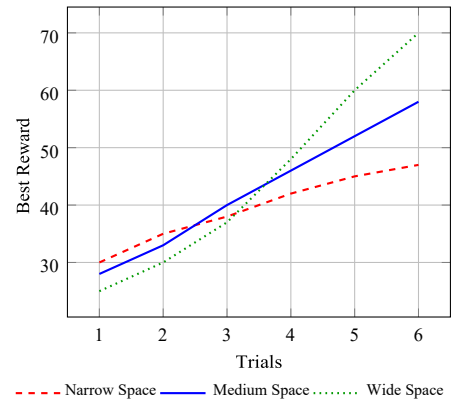


Fig.2. Convergence of Hyperopt on search spaces of size X, Y, and Z

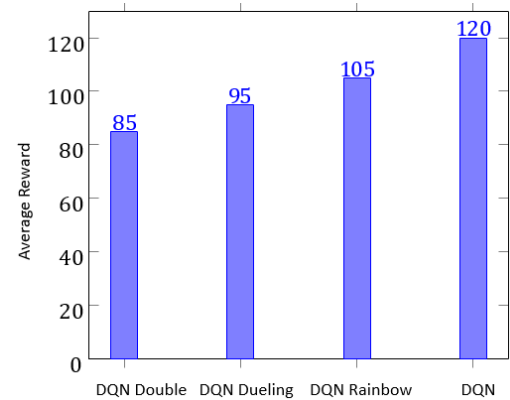


Fig.3. Comparison of performance of Rainbow DQN with other DQN algorithms

Our experiments showed that Bayesian optimization is a highly effective tool for navigating the complex hyperparameter space of Rainbow DQN. Optimized agents consistently outperformed baseline configurations, demonstrating improved training stability, faster convergence, and higher final rewards. Furthermore, comparisons with classical DQN variants affirmed the superiority of Rainbow DQN when augmented with carefully tuned hyperparameters. Ablation studies provided deeper insight into the contribution of individual architectural components within Rainbow DQN.

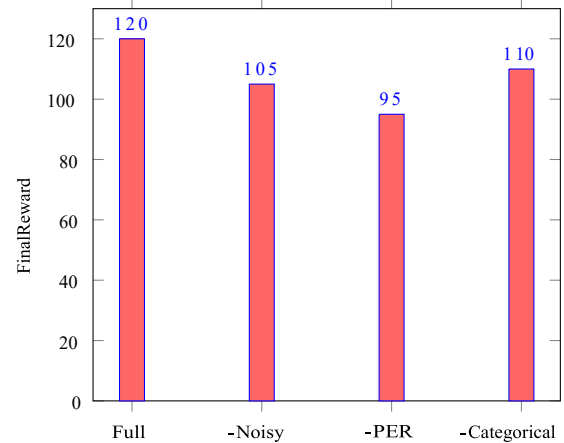


Fig.4. Effect of removing individual components of Rainbow DQN on performance.

7. CONCLUSION

Notably, mechanisms such as Prioritized Experience Replay and Noisy Networks emerged as critical to performance, underscoring the value of their inclusion in the integrated framework. Overall, this study highlights the practical value of combining advanced reinforcement learning architectures with principled, data-efficient hyperparameter optimization methods. The findings serve as a guideline for researchers and practitioners aiming to deploy high-performing agents with minimal manual tuning effort. Future directions may include extending this approach to more complex continuous control environments, integrating multi-objective optimization, or exploring meta-learning techniques to further enhance sample efficiency and generalization across tasks.

8. FUTURE WORK

While this study demonstrates the potential of Bayesian optimization for fine-tuning Rainbow DQN hyperparameters in a discrete control task, several promising avenues remain for future exploration.

- First, the methodology can be extended to more complex and high-dimensional environments, such as continuous control tasks in MuJoCo or robotic domains. These environments pose additional challenges in terms of sample efficiency and generalization, providing a robust testbed for evaluating the scalability of our optimization approach.
- Second, instead of using static search spaces, future research could adopt adaptive or hierarchical search strategies where the bounds of the hyperparameter space evolve based on previous performance. This dynamic adjustment could further improve convergence rates in large-scale or resource-constrained scenarios.
- Third, incorporating multi-objective optimization could enable simultaneous tuning for multiple goals—such as maximizing performance while minimizing training time or memory consumption. This would be especially relevant in real-world deployments where trade-offs between efficiency and accuracy are critical.
- Additionally, integrating meta-learning or AutoRL frameworks could allow agents to generalize tuning knowledge across different environments. Such approaches may reduce the need to repeat expensive tuning for every new task, significantly improving overall deployment efficiency.
- Lastly, while this work focused on the TPE algorithm, exploring alternative Bayesian strategies such as Gaussian Processes or Tree-based Sequential Model Optimization (SMBO) methods may yield better performance under certain constraints or domains.

By addressing these directions, future work can make hyperparameter optimization in deep reinforcement learning more scalable, generalizable, and applicable to real-world intelligent systems.

REFERENCES

- [1] M. Hessel, M. Joseph, H.V. Hassel, S. Tom, D. Will and H. Dan, “Rainbow: Combining Improvements in Deep Reinforcement Learning”, *Artificial Intelligence*, pp. 1-14, 2018.
- [2] J. Snoek, H. Larochelle and R.P. Adams, “Practical Bayesian Optimization of Machine Learning Algorithms”, *Machine Learning*, Vol. 45, pp. 1-12, 2012.
- [3] S.B. James, B. Remi, Y. Bengio and B. Kegl, “Algorithms for Hyper-Parameter Optimization”, *Proceedings of International Conference on Advances in Neural Information Processing Systems*, pp. 1-9, 2011.
- [4] M. Munir, S.W. Zamir and N. Memon, “Hyperparameter Optimization in Deep Reinforcement Learning: A Survey”, *Journal of Student Research*, Vol. 14, No. 1, pp. 1-9, 2025.
- [5] B. Shahriari, K. Swersky, Z. Wang, R.P. Adams and N. De Freitas, “Taking the Human Out of the Loop: A Review of Bayesian Optimization”, *Proceedings of the IEEE*, Vol. 104, No. 1, pp. 148-175, 2016.
- [6] H. Rakotoarison, M. Schoenauer and M. Sebag, “Automated Reinforcement Learning: A Survey”, *Proceedings of International Conference on Artificial Intelligence*, pp. 1-7, 2025.
- [7] L. Li, K. Jamieson, G. DeSalvo, A. Rostamizadeh and A. Talwalkar, “Hyperband: A Novel Bandit-based Approach to Hyperparameter Optimization”, *Proceedings of International Conference on Learning Representations*, pp. 1-8, 2017.
- [8] J. Bergstra, D. Yamins, and D.D. Cox, “Making a Science of Model Search: Hyperparameter Optimization in Hundreds of Dimensions for Vision Architectures”, *Proceedings of International Conference on Machine Learning*, Vol. 28, No. 1, pp. 115-123, 2013.
- [9] T. Akiba, S. Sano, T. Yanase, T. Ohta and M. Koyama, “Optuna: A Next-Generation Hyperparameter Optimization Framework”, *Knowledge Discovery in Databases*, pp. 1-10, 2019.
- [10] S. Falkner, A. Klein and F. Hutter, “Bohb: Robust and Efficient Hyperparameter Optimization at Scale”, *Proceedings of International Conference on Machine Learning*, pp. 1-19, 2018.
- [11] M. Janner, Q. Li and S. Levine, “Evaluating and Improving Automated Hyperparameter Tuning in Reinforcement Learning”, *Proceedings of International Conference on Machine Learning*, pp. 1-11, 2021.
- [12] Y. Song, P. Wang, Q. Wang and Y. Yang, “Auto-Tuning Hyperparameters in Deep Reinforcement Learning with Bayesian Optimization”, *Neurocomputing*, Vol. 423, pp. 885-894, 2021.
- [13] V. Mnih, K. Koray, D. Silver, A.A. Rusu, V. Joel and G.B. Marc, “Human-Level Control through Deep Reinforcement Learning”, *Nature*, Vol. 518, No. 7540, pp. 529-533, 2015.
- [14] H. Van Hasselt, A. Guez and D. Silver, “Deep Reinforcement Learning with Double Q-Learning”, *Proceedings of International Conference on Artificial Intelligence*, pp. 1-13, 2016.
- [15] Z. Wang, T. Schaul, M. Hessel, H. Van Hasselt, M. Lanctot and N. De Freitas, “Dueling Network Architectures for Deep

- Reinforcement Learning”, *Proceedings of International Conference on Machine Learning*, pp. 1-15, 2016.
- [16] T. Schaul, J. Quan, I. Antonoglou and D. Silver, “Prioritized Experience Replay”, *Proceedings of International Conference on Learning Representations*, pp. 1-21, 2016.
- [17] M. Fortunato, M.G. Azar, B. Piot, M. Jacob, O. Ian, G. Alex and M. Remi, “Noisy Networks for Exploration”, *Proceedings of International Conference on Learning Representations*, pp. 1-21, 2018.
- [18] M.G. Bellemare, W. Dabney and R. Munos, “A Distributional Perspective on Reinforcement Learning”, *Proceedings of International Conference on Machine Learning*, pp. 1-19, 2017.
- [19] Y. Burda, H. Edwards, A. Storkey and O. Klimov, “Exploration by Random Network Distillation”, *Proceedings of International Conference on Learning Representations*, pp. 1-17, 2019.
- [20] T. Xu, Q. Liu, L. Zhao and J. Peng, “Learning to Explore with Meta-Policy Gradient”, *Proceedings of International Conference on Machine Learning*, pp. 1-10, 2018.
- [21] P. Georgios, C. Flippis, S. Lukas and S.V. Albercht, “Benchmarking Multi-Agent Deep Reinforcement Learning Algorithms”, *Proceedings of International Conference on Advances in Neural Information Processing Systems*, pp. 1-9, 2021.
- [22] M. Jacob, O. Ian, G. Alex and M. Remi, “Learning to Optimize Hyperparameters Online During Training”, *Proceedings of International Conference on Learning Representations*, pp. 1-7, 2019.
- [23] A.A. Rusu, V. Joel and G.B. Marc, “Discovering States and Actions in Reinforcement Learning using Hierarchical Generative Models”, *Proceedings of International Conference on Machine Learning*, pp. 1-11, 2020.
- [24] Y. Wu, R.R. Afshar, Z. Rui and S. Wen, “Auro1: From Non-Parameterized to Parameterized Reinforcement Learning”, *Proceedings of International Conference on Artificial Intelligence*, pp. 1-19, 2020.
- [25] F. Hutter, L. Kotthoff and J. Vanschoren, “*Automated Machine Learning: Methods, Systems, Challenges*”, Springer, 2019.
- [26] T. Elsken, J.H. Metzen and F. Hutter, “Neural Architecture Search: A Survey”, *Journal of Machine Learning Research*, Vol. 20, No. 55, pp. 1-21, 2019.
- [27] B. Greg, C. Vicki, S. Jonas and S. John, “OpenAI Gym” *Proceedings of International Conference on Machine Learning*, pp. 1-5, 2016.
- [28] H. Peter, R. Islam, B. Philip, J. Pineau, P. Doina and D. Meger, “Deep Reinforcement Learning that Matters”, *Proceedings of International Conference on Artificial Intelligence*, pp. 1-26, 2018.
- [29] M. Hessel, J. Modayil, H. Van Hasselt, T. Schaul, G. Ostrovski, W. Dabney, D. Horgan, B. Piot, M.G. Azar and D. Silver, “Rainbow: Combining Improvements in Deep Reinforcement Learning”, *Computing Research Repository*, Vol. 3, pp. 1-14, 2017.