REST Journal on Data Analytics and Artificial Intelligence
Vol: 4(1), March 2025
REST Publisher; ISSN: 2583-5564
Website: http://restpublisher.com/journals/jdaai/
DOI: https://doi.org/10.46632/jdaai/4/1/3

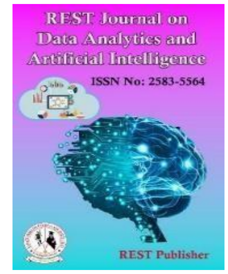

Reinforcement Learning for Autonomous Driving

A. Manikanta, B. Harsha vardhan, P. Suchith, K. BasavaRaju
School of Engineering Anurag University, Hyderabad, India.
**Corresponding Author Email: manikantaakula04@gmail.com*

Abstract. *Reinforcement Learning (RL) has emerged as a powerful approach for training autonomous driving systems by enabling vehicles to learn optimal driving policies through interaction with their environment. Unlike rule-based or supervised learning methods, RL allows agents to adapt to dynamic and complex traffic scenarios without explicit programming. This paper explores various RL techniques, including Deep Q-Networks (DQN), Policy Gradient Methods, and Model-Based RL, applied to autonomous driving tasks such as lane keeping, obstacle avoidance, and decision-making at intersections. Key challenges, including sample efficiency, safety constraints, and real-world generalization, are discussed along with potential solutions such as hybrid learning models and simulation-based training. The study highlights recent advancements in RL for autonomous driving and outlines future research directions to enhance robustness and deployment in real-world applications.*

Keywords*: Reinforcement Learning (RL),Autonomous Driving, Deep Q-Networks (DQN),Policy Gradient Methods, Decision-Making, Sensor Fusion, Simulation-based Training, Reward Function, Safe Exploration, Real-World Deployment.*

1. INTRODUCTION

Autonomous driving is a complex task requiring real-time decision-making in dynamic and unpredictable environments. Traditional rule-based and supervised learning [1-2] approaches often struggle to handle the variability of real-world traffic scenarios. Reinforcement Learning (RL) has emerged as a promising solution, enabling autonomous vehicles to learn optimal driving policies through trial and error, maximizing long-term rewards while ensuring safety and efficiency.RL frameworks, such as Deep Q-Networks (DQN) and Policy Gradient methods, allow agents to learn from simulated and real-world interactions [2-5]. These methods help vehicles make decisions in various driving tasks, including lane following, obstacle avoidance, and intersection navigation. However, challenges such as sample inefficiency, safety constraints, and generalization to real-world environments remain key hurdles.This paper explores RL techniques in autonomous driving, their advantages over traditional approaches, and the ongoing research to improve safety, robustness, and deployment feasibility. function” Correlations have been made between animal behavior and reinforcement learning [6-7],

Reinforcement Learning is a part of machine learning paradigms along with Supervised and unsupervised learning, focused on a cumulative award system by which an agent accomplishes a certain task/goal. The purpose of reinforcement learning is for the agent to learn an optimal or nearly-optimal policy that maximizes the “reward. function” Correlations have been made between animal behavior and reinforcement learning,for example pleasure and food is seen as a positive reinforcement and animals engage in behaviors that optimize these rewards. Machine reinforcement learning works in the same manner. It tries to maximize the amount of reward it gets by checking the different possibilities/routes it can take. Reinforcement learning requires clever exploration mechanisms and it ran domly selects actions without reference. Instead of learning from a data-set, it learns from mistakes generated by a reward system, so our robots basically learn from experience and mistakes. If an action is positive, it will have a big reward, however if the action is negative then the reward is negative. The goal of reinforcement learning is to have the agent choose the action that gets the biggest reward [8]. Self-driving cars are built with a specific algorithm and are not made smart. The algorithms that govern and control their movements are continuously learning as the technology is being prominent in cars such as Tesla. Self-driving cars are one of the many examples of the power of machine learning in the technological world [9-13]. Autonomous driving systems are multiple perception level tasks which have now achieved very high results using deep learning architectures. Apart from perception autonomous driving systems are multiple tasks where the more classic supervised learning methods are nomoreapplicable. Machine learning is crucial in the autonomous car industry. It has already been applied in some aspects of vehicles, such as advanced driver assistance systems. One of the areas in which autonomous cars struggle today is the classification of objects on the road. This is the most crucial concept in autonomous vehicles because they must classify if a sign is a stop sign or the differences between a pole or a person. There is no room for error in these situations because of how dangerous being on the road can be; any small mistake can be the difference between life or death. Reinforcement learning plays a significant role in this area because the more realistic our programs are, the better the machines will perform, hence why virtual reality is starting to play a substantial role in autonomous vehicles [14]. Traffic simulators are important ways companies make improvements to their auto-driving systems; they give real-life scenarios without the dangers that come with them. Autonomous learning in cars consists of a reward system; if a car performs well, it gets rewarded a point, and if it fails a particular task, such as moving out of its lane, a point gets deducted [15].

One notable project done on Reinforcement Learning was at Nanyang Technological University and focused on designing and implementing a neural network that maximises driving speed of self-driving cars through reinforcement learning. The project implements reinforcement learning to generate a self-driving car-agent with a deep learning network to maximize its speed. The convectional neural network was implemented to extract features from a matrix representing the environment mapping of self-driving cars. The model acts as value functions for five actions estimating future rewards. The model is trained under Q-learning algorithm in a simulation built to simulate traffic conditions of a seven-lane expressway. After continuous training for 2340 minutes, the model learns the control policies for different traffic conditions and reaches an average speed 94 km/h compared to maximum speed of 110 km/h [16]

2. LITERATURE SURVEY

Reinforcement Learning (RL) has emerged as a powerful framework for autonomous driving, enabling vehicles to learn complex decision-making strategies through and supervised learning approaches often struggle to generalize across diverse driving scenarios, whereas RL methods offer adaptability and robustness in dynamic and uncertain conditions [17-18]. Researchers have explored various RL techniques, including Deep Q-Networks (DQN), Policy Gradient methods, and Actor-Critic architectures, to enhance the learning efficiency and decision-making capabilities of autonomous vehicles.

Several studies have demonstrated the effectiveness of RL in solving tasks such as lane keeping, obstacle avoidance, and traffic signal management. For instance, Deep Q-Networks have been successfully applied to learn optimal policies for navigation in simulated urban environments, achieving near-human performance. Additionally, policy-based methods like Proximal Policy Optimization (PPO) and Trust Region Policy Optimization (TRPO) have shown promise in handling continuous action spaces, allowing for smoother and safer driving maneuvers. Multi-agent RL frameworks have also been proposed to enable coordination among multiple autonomous vehicles, improving traffic efficiency and reducing congestion [19].

Despite significant progress, RL-based autonomous driving faces several challenges, including sample inefficiency, safety concerns, and the reality gap between simulation and real-world deployment. Training RL models requires a large amount of interaction data, which is often impractical to collect in real-world settings. To address this, researchers have explored model-based RL and imitation learning techniques to reduce the need for extensive training data. Furthermore, safety-critical applications demand robust RL policies that ensure collision avoidance and regulatory compliance. Hybrid approaches combining RL with rule-based safety constraints and heuristic planning have been investigated to enhance the reliability of autonomous driving systems [20].

Study	Challenges Addressed
Mnih et al. (2015)	Sample inefficiency, generalization
Lillicrap et al. (2016)	Stability in training
Schulman et al. (2017)	Semple efficience,
Liang et al. (2020)	Coordination challenges

Reinforcement Learning (RL) has gained significant attention in autonomous driving due to its ability to learn optimal driving policies through interactions with the environment. Unlike supervised learning, which relies on labeled data, RL enables vehicles to make real-time decisions by maximizing cumulative rewards. Deep Reinforcement Learning (DRL), which integrates deep learning with RL techniques, has further improved decision-making in complex traffic environments. Researchers have applied various RL algorithms such as Deep Q-Networks (DQN), Proximal Policy Optimization (PPO), and Soft Actor-Critic (SAC) to achieve better lane-keeping, adaptive cruise control, and collision avoidance.

A key challenge in RL for autonomous driving is ensuring safety and sample efficiency. Since real-world data collection is expensive and risky, researchers rely on high-fidelity simulation platforms like CARLA, SUMO, and Air Sim to train and test RL models. Sim-to-real transfer learning methods, such as domain adaptation and imitation learning, are used to bridge the gap between virtual and real-world environments. Hybrid approaches that combine RL with classical control methods, including Model Predictive Control (MPC) and heuristic rule-based techniques, have been explored to improve the stability and interpretability of RL policies. Moreover, multi-agent RL frameworks have been studied for cooperative driving, enabling multiple autonomous vehicles to interact and optimize traffic flow.

Recent studies have focused on improving the robustness of RL-based autonomous systems under uncertain and dynamic conditions. Adversarial learning techniques have been used to test the resilience of RL models against unexpected scenarios, such as sudden pedestrian crossings or aggressive driving behavior. Additionally, hierarchical RL approaches, where high-level policies guide low-level control actions, have shown promise in managing complex urban driving tasks. Future research directions include enhancing the generalization capabilities of RL models, integrating RL with advanced perception systems, and developing explainable AI techniques to improve trust and transparency in autonomous driving decisions.

Recent advancements in simulation environments, such as CARLA and SUMO, have facilitated large-scale RL experiments, bridging the gap between research and deployment. Transfer learning and domain adaptation techniques are being actively studied to improve the generalization of RL policies across different driving conditions. As RL algorithms continue to evolve, integrating them with sensor fusion, computer vision, and V2X (vehicle-to-everything) communication is expected to further enhance autonomous driving capabilities. Future research directions include improving sample efficiency, ensuring interpretability of RL decisions, and developing more robust real-world testing frameworks to accelerate the adoption of RL in autonomous driving.

3. DATASET

Dataset Description: Machine learning is crucial in the autonomous car industry. It has already been applied in some aspects of vehicles, such as advanced driver assistance systems. One of the areas in which autonomous cars struggle today is the classification of objects on the road. This is the most crucial concept in autonomous vehicles because they must classify if a sign isa stop sign or the differences between a pole or a person. There is no room for error in these situations because of how dangerous being on the road can be; any small mistake can be the difference between life or death. Reinforcement learning plays a significant role in this area because the more realistic our programs are, the better the machines will perform, hence why virtual reality is starting to play a substantial role in autonomous vehicles. Traffic simulators are important ways companies make improvements to their auto-driving system hey give real-life scenarios without the dangers that come with them. Autonomous learning in cars consists of a rewards system; if a car performs well, it gets rewarded a point, and if it fails a particular task, such as moving out of its lane, a point gets deducted. A recent project done by Stanford University gives people a good understanding of how these systems work. Here is a look at the reward algorithm function that was used in their project [15]. This is a lower-end system, and the caveat to this is that it only penalizes the machine if it does the opposite of what it was told to do. In more serious cases such as developments in Tesla, the reward system is much more diverse and has a lot more depth to it, but this gives people a basic understanding of how it works.

4. FINDINGS

Surveys Conducted Moving forward, evaluating the effectiveness of self-driving algorithms was an important step. A study done by Christian Berger et al. presents a survey for publicly available data sets next to an overview of virtual testing environments to support the research, development, and evaluation of algorithms from the field of autonomous driving. They presented 37 data sets from different perspectives such as included driving situations, sensor setups, data format; additionally, up to 22 virtual testing environments are presented to support the closed-loop testing. This work is the first and most comprehensive survey of this kind providing guidance about existing and publicly available assets to researchers and developers [10]. The goal of this work was to carefully conduct a broad survey to provide an exhaustive overview about existing datasets and virtual testing environments supporting research and development of autonomous driving and algorithms. This work focuses particularly on data sets and virtual testing environments that could be identified using structured web searches and systematic snowballing and that are accessible to researchers and developers. For the data sets, they set their focus only on ground truth driving data collected on public roads with partial or full open access. For the virtual testing environments, they focused especially on solutions available as open source to encourage and facilitate contributions from Algorithm 1 Reward Function procedure REWARD (st,at)Vmin ←5 Defined in meters/second. Brake ←20 Defined in meters/second.Vmax ←5 Defined in meters/second. Threshold ←0.05 Defined in meters/second. Reward ←0if Vt> Vbrake andacct=Brake thenreward ←reward −1end ifif Vt> Vbrake andacct=Accelerate thenreward ←reward + 1end ifif Vt¡ Vmax and acct= Accelerate then reward ←reward + 1end ifif Vt¡ Vmax and acct=Brake thenreward ←reward −1end ifif φt<−φthreshold andSteert=T urnRig ht thenreward ←reward + 1end ifif φt<−φthreshold and Steert=T urnLef t thenreward ←reward −1end ifif φt> threshold and Steert=T urnLef t then reward ←reward + 1end ifif φt> φthreshold andSteert=T urnRig ht thenreward ←reward −1end ifif Vt< Vminandacct6=Accelerate thenreward ←reward −1end if return reward end procedure the community. A deep analysis per data-set or virtual testing environment is very specific to particular use-cases of the development or evaluation of autonomous driving. Hence it would not contribute to the overall goal of the work but is rather suggested in specialized subsequent studies [9]. B. Potential Obstacles Along with the proper usage of these self-learning au-tumorous cars, a couple of issues have been identified. The main addressed are issues related to power usage, effectiveness of sensors being used along with the accuracy of communication links and technologies and safety of drivers. Most of these vehicles are designed to address the air pollution and reduce greenhouse gas emissions, and develop a higher fuelefficiency. So far, most of the existing energy management system strategies are either just simply following predefinedrules that are not adaptive to changing driving conditions; or heavily relying on accurate prediction of future traffic conditions. Deep reinforcement learning does not have a system

designed to adapt to conditions that would be a mutafacient use of energy. Progress has been made in the past on how the car manufacturing industry tries to implement ineffective usage of power. For example, Tesla uses thousands of lithium-ion cells to power the vehicle. Equipped with a heating system capable of warming the battery in cool temperatures, Tesla motors are as durable as they are capable. Unlike hybrid motors, Tesla's batteries have to be charged from outlets. Thesis because they run solely on battery power, meaning the electric motor cannot be charged automatically while driving [2]. C. Comparison of Reinforcement Learning techniques used by automotive companies Model-based reinforcement learning strongly derives from the control theory. The goal is to plan through an f (s, a) control function and choose the most optimal actions. Model-based learning uses a Reinforcement learning field where the creator gives the laws of physics; this allows for many different simulations to be run. One drawback that comes with this time of RL is that they have more assumptions and approximations when given a task. Within Model-Based RL, there are two main approaches; learning the model or learn given the model. The learn the model is base policy ran; the trajectory is observed just like a random or educated policy. Then, the model is fitted using the sample data that has been collected. Another method is the supervised learning tactic used to Traina model to minimize the errors from sampled data. There is also a cost function used in this method to measure how far they are from the target location and how much effort is spent. There are also the Imagination-Augmented Agents that learn to interpret predictions based on a learned environment model. They then construct implicit plans in arbitrary ways and use those predictions as an additional context within deep policy networks. This is considered a hybrid learning method because it incorporates some features of model-based and model-free learning methods on the contrary, there are Model-Free learning methods that are also used in reinforcement learning techniques. The two main approaches used in Model-Free learning are policy optimization and Q-Learning. In Policy optimization, the agent directly learns the policy function that maps state to action, determined without a value function. It uses an equation that finds the best parameters to optimize a score function; the discount factor and the reward are also factored in the equation. The main steps are to ensure the quality of the policy score function and find the correct gradient ascent to provide the best parameters that improve policy over time. On the other hand, Q-Learning learns the action-value function, which determines how good it would be to perform aseptic action at a particular state. A scalar value is assigned over an action given the form; the typical algorithm starts by choosing an action. Then, it acts, measures the rewards, and finally updates the Q. Q-tables can be very complex when working on big space projects, so instead, a Deep Neural network is used. These neural networks approximate the Q-values for each action based on the current state. Many other methods are used in reinforcement learning, but these are some of the more contemporary and feasible scenarios used in today's industries. Each state in the environment will be represented by a pixel and the agent may be able to take several actions from each state. The iterative process of acquiring and updating q values for each action pair in a large space like the case of data from the car vision becomes computationally inefficient. Rather than using value iterations to directly compute Q values and find the optimal Q function, we instead use a function approximator to estimate the optimal function and this is done by the artificial neural networks. The deep neural network is used for each state action pairing a given environment and in turn network will approximate the optimal Q function. A deep Q algorithm tries to find a maximum value for the q function given below $q(s, a) = E\ [Rat\ (t+1) + \gamma max\ q(s1, a1)$ where s is the state of the agent, a is the action by the agent, Rat (t+1) is the reward that the agent gets after a time step of t+1, and gamma is the maximum value of q for the inputs s prime and a prime [14].Many of the companies that implement autonomous driving systems implement these reinforcement learning methods dis-cussed above. For example, Waymo, Tesla, Ford, and GM all implement some methods of RL. This portion will focus on how these companies implement strategies and how they are similar to one another. Waymo, whose parent company is Google, uses deep Mind and deep reinforcement learning to create agents and drive-in policies. Policies are basic behaviors, and Waymo can simulate angry drivers cutting people off, careless bikers. Their programs evaluate the algorithms to ensure that they are corrected to the best of their abilities. After they have accurate and trained predictions on the drivers, they generate the best trajectories to take, which is called decision-making and trajectory generation, which for Waymo, is called Chauffeur Net. [6] The goals of the trajectory models are to set way points that have the lowest errors in terms of safety, speed, and feasibility. Lastly, the trajectory considers attractors and resellers. This helps cars stay in the center lane and avoid barricades. Tesla uses similar techniques as Waymo; however, Tesla claims to drive nearly 15 million miles per day, which is almost the amount Waymo has driven in all of its existence. Tesla uses computer vision, prediction, and path planning to aid its reinforcement learning methods. they use neural networks to annotate images so that the computer can correctly determine how to handle the situation. They use thousands of pictures that have been correctly annotated to help the machine accurately distinguish between different objects. The more pictures that are used, the better the machine can get at learning these objects. Tesla also uses the neural network to imitate how humans interact, which is known as imitation learning. In imitation learning, the neural networks learn to predict how human drivers would do by drawing a correlation between what the computer sees and how the humans took action. The neural network helps Tesla in most of its reinforcement learning techniques and is a big reason why they are considered the most advanced autonomous driving company

Ford is another company that uses artificial intelligence and machine learning techniques in its autonomous driving advancements. Ford also uses the neural network to aid its company; however, they use it for quality assurance, such as supporting their supply change and detecting malfunctions. Ford also used machine learning to provide features, personal-ization, and product improvements to the consumer. Much of these features interact directly with the drivers. For example, they aid in acceleration/braking, improving fuel efficiency, and climate control. However, Ford does not implement machine learning in the same ways that Tesla and Waymo do. Ford seems to implement it more in correlation with the drivers rather than the driving technology. Machine learning plays a huge role in autonomous driving at Ford; however, it doesn't seem as apparent as the other two large companies [12].GM is another company that implements machine learning in its autonomous vehicle research. General Motors Cruise is a company that focuses on developing fully autonomous robotaxis. GM has one of the most reliable and sophisticated machines learning prediction systems. Their systems give self-driving prototypes knowledge to read the road ahead, allowing them to predict better what other cars or pedestrians may do. The goal of the GM system is to excel within long-tail events. Long-tail events are considered to be rare events such as U-turns or pedestrians stepping in front of vehicles. Engineer slog miles on their prototypes and begin to log the infrequent outlier events; they then use up-sampling techniques to teach the machine learning systems how to handle these rare events.GM also excels in data labeling, which is their ability to familiarize command items on the roads. For example, when they correctly label a trajectory, it records their memory and aids in its self-learning functions. GM is one of the top leaders in machine learning, and they look to continue their efforts in their autonomous taxi company [11]. D. Result As mentioned earlier, the Q learning system is a simple and effective system for a system with a small input. However, when the input involves a higher dimensional sensory, then we start to see issues and end up looking for new time and 'storage' effective methods to do so. Despite the effort that is put in, there is a lot of windows available to install a new system. When considering the levels of Driving Automation, currently, we are in a Level two partial driving automation state. This is a stage where the vehicle consists of an advanced driver assistance system where it can control both the steering and accelerating/decelerating. However, we expect to see more changes with the following levels of vehicle autonomy. Level three autonomy, which consists of monitoring the driving environment. In this case, the vehicle is able to analyze what is going on in the environment and make informed decision after that like passing by a vehicle moving slow. Audi is the only company that has been working on this and is in the process of releasing their level three A8L autonomous vehicle in the Germany in the near future. The next one is a level four Automation (High Driving) which is similar to the level three automation, with an addition of an interaction with the environment and response if there is any change in the system or if things go wrong. They don't usually require human assistance. The last stage is a level five where the autonomous vehicle doesn't contain an accelerator or a steering wheel that is available for the driver (technically a rider) [1]. Production of these cars requires an advanced method of deep learning. One major drawback is the fact that the deep learning algorithm doesn't work in different types of environments, for example even for a level four automation, it requires an urban environment where speed reaches an average of 30 mph). One other factor could also be the lack of data to be stored to maintain an autonomous vehicle. A current estimate suggests a level-five autonomous vehicle generates between 1 and 20 terabytes per hour during the development stage. However, we still haven't seen a vehicle of such abilities yet. An autonomous vehicle usually consists of Cameras with 2D and 3D images and video inputs, a Radar that uses long wavelength radio waves, Long and medium Range Radar, Ultrasonic sensor, and a Light Detection and Ranging. Despite having these advanced features, the data collection system is not sufficient enough to enable vehicles to drive outside of an urban state or in the wildlife. The solution we believe would be efficient is the collection of data even after the production and release of autonomous vehicles and using a different environment for the test cars initially. For example, in the United States on-road testing is occurring in 17 cities and it is not sufficient enough

5. CONCLUSION

In this project, we have tried to analyse the stage that the deep reinforcement learning is at with the different approaches like the Deep Q learning, Greedy approach, and Q learning. We have found that Deep Q learning has been the mutafacient way that most of the leading companies have been using. One of our main data sources is research done by Matt Vitelli and Aran Nayebi from Stanford University. They constructed a deep Q network (DQN) agent that performs the tasks of an autonomous car and compared its efficiencies that of the different agents used these days other tandan. They have found out that the DQN agent showed a more stable behavior, maximum driving speed, and a higher cumulative reward as compared to the results from a Greedy Agent, Discrete Agent and a Hand-Crafted Controller.

REFERENCES

[1]. Synopsys-the 6 levels of vehicle autonomy explained, 2021.
[2]. Myenteric- how does a tesla battery work? a guide to electric car batteries, Jan 2021.
[3]. Smart Lab AI. Reinforcement learning algorithms-an intuitive overview 2019.
[4]. Yarrow Bouchard. Tesla's deep learning at scale: Using billions of milesto train neural networks, Dec 2019.
[5]. Paul Christianson. Billions of miles of data: The autonomous vehicle training conundrum.
[6]. Jeremy Cohen. How google's self-driving cars work, Aug 2020.
[7]. Abdur R Fayjie, Sabir Hossain, Doukhi Oualid, and Deok-Jin Lee.Driverless car: Autonomous driving using deep reinforcement learning in urban environment. In 2018 15th International Conference onUbiquitous Robots (UR), pages 896–901. IEEE, 2018.
[8]. Y. Kang, H. Yin, and C. Berger. Test your self-driving algorithm: An overview of publicly available driving datasets and virtual testing environments. IEEE Transactions on Intelligent Vehicles, 4(2):171–185,2019.
[9]. Yue Kang, Hang Yin, and Christian Berger. Test your self-drivingalgorithm: An overview of publicly available driving datasets and virtualtesting environments. IEEE Transactions on Intelligent Vehicles, PP:1–1,03 2019.
[10]. B Ravi Kiran, Ibrahim Sobh, Victor Talpaert, Patrick Mannion, Ahmad AAl Sallab, Senthil Yogamani, and Patrick P´erez. Deep reinforcement learning for autonomous driving: A survey. IEEE Transactions on Intelligent Transportation Systems, 2021.
[11]. Sam McEachern, Sam McEachern Sam loves to write, has a passion for auto racing, Sam McEachern, Sam loves to write, and has apassion for auto racing. How cruise uses machine learning to predict the unpredictable.
[12]. Peter Sonnenberg. Broad use of machine learning at ford.
[13]. Jack Stilgoe. Machine learning, social learning and the governance of self-driving cars. Social studies of science, 48(1):25–56, 2018.
[14]. Andre Violante. Simple reinforcement learning: Q-learning, Jul 2019.
[15]. Matt Vitelli and Aran Nayebi. Carma: A deep reinforcement learning approach to autonomous driving. Tech. Rep. Stanford University, 2016
[16]. Manoranjan Dash, Design of Finite Impulse Response Filters Using Evolutionary Techniques - An Efficient Computation, ICTACT Journal on Communication Technology, March 2020, Volume: 11, Issue: 01
[17]. Manoranjan Dash et al.," Effective Automated Medical Image Segmentation Using Hybrid Computational Intelligence Technique", Blockchain and IoT Based Smart Healthcare Systems, Bentham Science Publishers, Pp. 174-182,2024
[18]. Manoranjan Dash et al.," Detection of Psychological Stability Status Using Machine Learning Algorithms", International Conference on Intelligent Systems and Machine Learning, Springer Nature Switzerland, Pp.44-51, 2022.
[19]. Manoranjan Dash, N.D. Londhe, S. Ghosh, et al., "Hybrid Seeker Optimization Algorithm-based Accurate Image Clustering for Automatic Psoriasis Lesion Detection", Artificial Intelligence for Healthcare (Taylor & Francis), 2022, ISBN: 9781003241409
[20]. Manoranjan Dash, N.D. Londhe, S. Ghosh, et al., Swarm intelligence-based clustering technique for automated lesion detection and diagnosis of psoriasis, Computational Biology and Chemistry. 86 (2020) 107247.

