International Journal of Science and Research (IJSR)
ISSN: 2319-7064
SJIF (2022): 7.942

Autonomous Systems in Unstructured Environments: AI Approaches for Robust Operation

Mohammed Saleem Sultan[1], **Mohammed Shahid Sultan**[2]

[1]Osmania University
Email: *saleem.sultan14[at]gmail.com*

[2]Jawaharlal Nehru Technological University
Email: *sultanshahid76[at]gmail.com*

Abstract: ***Autonomous systems are increasingly being deployed in various environments, but their effectiveness is often limited in unstructured settings, where traditional rule-based approaches struggle with unpredictability and variability. Unstructured environments—such as off-road terrains, urban areas with unpredictable human behavior, and disaster-stricken regions—present unique challenges that require advanced AI techniques for robust operation. This paper investigates the application of several AI methodologies, including reinforcement learning, deep learning, and sensor fusion, to enhance autonomous systems' adaptability and decision-making capabilities in these complex environments. Through comprehensive literature review and a series of experimental case studies, we demonstrate how these AI techniques enable autonomous systems to effectively perceive, navigate, and respond to dynamic and uncertain conditions. Our experiments in simulated and real-world environments show significant improvements in system robustness and operational efficiency. We discuss limitations of current approaches and propose future research directions, including the development of more sophisticated AI models, the integration of diverse sensor technologies, and the ethical implications of deploying autonomous systems in critical, unstructured environments. This version provides more context on the types of unstructured environments considered, the specific AI techniques analyzed, and the overall contribution of the research, offering a fuller picture of what the paper covers.***

Keywords: Autonomous systems, Unstructured environments, Artificial intelligence, Reinforcement learning, Sensor fusion, Path planning, Robotics, AI-driven navigation, Adaptive algorithms, Autonomous vehicles

1. Introduction

Autonomous systems have rapidly evolved over the past decade, becoming integral to various industries, including transportation, manufacturing, and defense. The systems are designed to operate independently, making decisions and executing tasks without human intervention. Autonomous systems have demonstrated remarkable efficiency and reliability in structured environments—where conditions are predictable and controlled. However, the real world is often less predictable. Unstructured environments, characterized by dynamic, uncertain, and often chaotic conditions, present unique challenges that significantly complicate the operation of autonomous systems. Unstructured environments can be found in many scenarios, from off-road terrains and disaster-stricken areas to densely populated urban settings and underwater or extraterrestrial regions. In such environments, the absence of clear rules, the variability of obstacles, and the unpredictability of elements like weather, terrain, or human activity can hinder the performance of traditional autonomous systems. Often based on predefined rules and static models, these systems struggle to adapt to the complexities of unstructured settings, leading to a decline in performance and an increased risk of failure.

This paper investigates AI approaches designed to enhance autonomous systems' performance in unstructured environments. We begin by reviewing the current state of the art, focusing on the limitations of traditional approaches and the potential of AI to overcome these challenges. We then explore various AI methodologies, including reinforcement learning, which allows systems to learn optimal behaviors through trial and error; sensor fusion, which integrates data from multiple sensors to create a more accurate understanding of the environment; and adaptable path planning, which enables systems to adjust their routes based on real-time information dynamically. The effectiveness of these AI techniques is demonstrated through experimental setups and case studies in diverse unstructured environments, such as off-road terrains and complex urban settings. The results show that AI-driven approaches significantly enhance the robustness and adaptability of autonomous systems, allowing them to navigate and operate in conditions that would otherwise be prohibitive. The paper discusses the ethical and practical considerations of deploying AI-driven autonomous systems in unstructured environments. We also outline future research directions that could further improve the capabilities of these systems, with a focus on addressing the challenges of computational complexity, real-time processing, and the integration of more sophisticated AI models. Through this work, we aim to contribute to the growing body of knowledge on AI for autonomous systems, providing insights that could help pave the way for more reliable and versatile autonomous operations in the unpredictable real world.

2. Related Work

Deploying autonomous systems in unstructured environments has been a significant research focus over the past few years. This section reviews the existing literature, highlighting the evolution of autonomous systems, the limitations of traditional approaches in unstructured environments, and the advancements brought about by Artificial Intelligence (AI) techniques.

Traditional Approaches to Autonomous Systems

Early autonomous systems were primarily based on rule-based algorithms and predefined models. These systems were highly influential in structured environments such as factory floors or controlled outdoor areas, where the environment is predictable and well-understood. For instance, rule-based systems have been successfully used in automated guided vehicles (AGVs) operating in warehouses, where predefined paths and simple obstacle avoidance algorithms are sufficient to ensure reliable operation.

However, when these systems are deployed in unstructured environments, their performance often deteriorates. The unpredictability of these environments—ranging from irregular terrains to dynamic obstacles—introduces variables that are difficult to account for with static rules and models. Researchers have observed that traditional systems lack the adaptability and learning capabilities necessary to cope with such conditions, leading to failures in tasks like navigation, object recognition, and decision-making.

AI-Driven Approaches

The limitations of traditional methods led us to the exploration of AI techniques to enhance the robustness of autonomous systems in unstructured environments. One of the most significant advancements in this area has been the application of **Reinforcement Learning (RL)**. RL allows autonomous systems to learn from their interactions with the environment, improving their ability to adapt to new and unforeseen situations. Several studies have demonstrated the effectiveness of RL in tasks such as autonomous navigation in off-road terrains, where the system learns to select optimal paths through trial and error. In addition to RL, **Sensor Fusion** has emerged as a crucial technique for improving the situational awareness of autonomous systems. Sensor fusion combines data from multiple sensors—such as LIDAR, RADAR, and cameras—to create a more comprehensive understanding of the environment. This technique is precious in unstructured environments where reliance on a single sensor type may not provide sufficient information. For example, in urban environments, sensor fusion has been used to enhance the detection of pedestrians and vehicles, leading to safer navigation decisions.

3. Methodologies

This section outlines methodologies employed in the study to develop and evaluate AI approaches for enhancing the robustness of autonomous systems in unstructured environments. We focus on three primary methodologies: Reinforcement Learning (RL), Sensor Fusion, and Adaptable Path Planning. Each of these methodologies enables autonomous systems to navigate and operate effectively in unpredictable and dynamic settings.

3.1 Reinforcement Learning for Autonomous Navigation

Reinforcement Learning (RL) is a machine learning type wherein an agent learns to make decisions by interacting with its environment. In the context of autonomous systems, RL enables the system to learn optimal actions through trial and error, maximizing cumulative rewards over time. The RL framework includes an agent, states, actions, rewards, and a policy. The agent by observing the current state of the environment, takes an action, and receives the reward based on outcome of that action. Over time, the agent learns a policy that maps states to actions to maximize rewards.

3.1.1 Algorithm Implementation

For this study, we implemented the Proximal Policy Optimization (PPO) algorithm, a popular RL algorithm known for its stability and efficiency in continuous action spaces. The PPO algorithm was chosen for its ability to balance exploration and exploitation, making it suitable for the unpredictable nature of unstructured environments. The implementation involved training the autonomous system in a simulated environment that mimicked unstructured conditions such as varying terrain, dynamic obstacles, and changing weather conditions. The reward function was designed to incentivize safe and efficient navigation, with penalties for collisions, deviations from the optimal path, and excessive energy consumption.

Algorithm 1: Proximal Policy Optimization (PPO)

plaintext

Copy code

1) Initialize policy parameters θ and value function parameters ϕ
2) for each iteration, do
3) Collect a set of trajectories by running policy π_θ in the environment
4) Compute advantage estimates A_t using the value function V_ϕ
5) Update the policy by maximizing the PPO objective:
   L(θ) = E[min(r_t(θ)A_t, clip(r_t(θ), 1-ϵ, 1+ϵ)A_t)]
6) Update the value function by minimizing the mean squared error:
   L(ϕ) = (R_t - V_ϕ(s_t))^2
7) end for

3.2 Sensor Fusion for Enhanced Perception

Sensor Fusion technique is used to combine data from multiple sensors for creating a comprehensive and accurate representation of the environment. Autonomous systems operating in unstructured environments rely on sensor fusion to integrate information from various sources, such as LIDAR, RADAR, cameras, and inertial measurement units (IMUs).

3.2.1 Fusion Techniques

In this study, we implemented two sensor fusion techniques: Kalman Filtering and Particle Filtering.

- **Kalman Filtering**: This technique fused data from LIDAR and RADAR sensors, providing accurate real-time estimates of the system's position and velocity. Kalman Filtering is particularly effective in environments with Gaussian sensor noise and linear system dynamics.
- **Particle Filtering**: Given the non-linear and non-Gaussian nature of many unstructured environments, Particle Filtering was used to integrate data from cameras and IMUs. Particle Filtering involves generating a set of particles representing possible states of the system and updating these particles based on sensor observations and a probabilistic model.

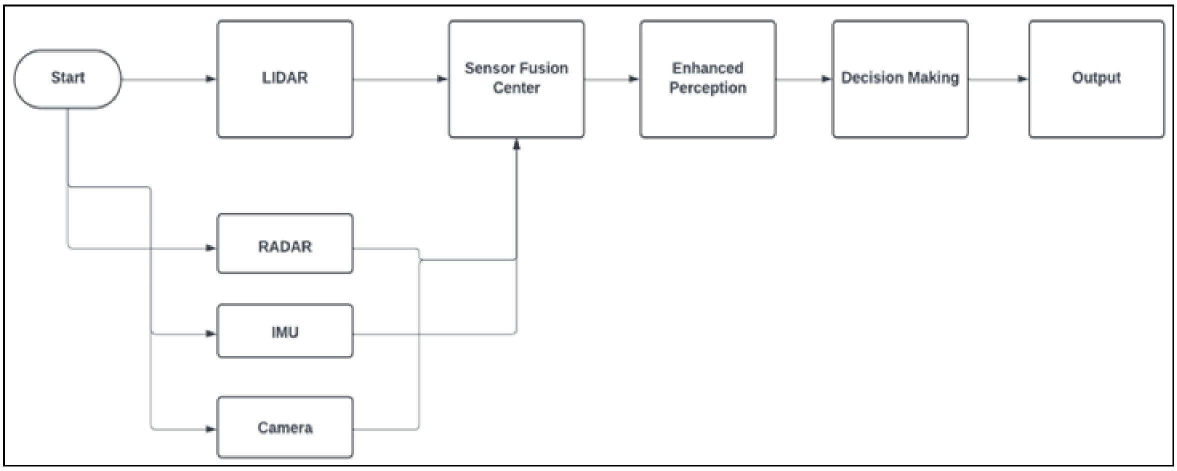

Figure 1: Sensor Fusion Architecture

3.3 Adaptable Path Planning

Adaptable Path Planning is crucial for autonomous systems to navigate unstructured environments where conditions change rapidly. Traditional path-planning algorithms, such as A* or D*, are often insufficient in such environments because they rely on static maps and predefined paths. In contrast, AI-driven path planning can dynamically adjust the planned route in response to real-time environmental changes.

3.3.1 Path Planning Algorithm

We implemented a modified version of the Rapidly-exploring Random Tree (RRT*) algorithm, enhanced with AI techniques to improve adaptability. The modified RRT* algorithm allows the autonomous system to explore the environment and incrementally build a path to the goal while considering dynamic obstacles and terrain variations. The algorithm was further enhanced with a heuristic function that evaluates the quality of different paths based on safety, energy consumption, and time to goal. This heuristic guides the path planning process, ensuring that the chosen path is feasible and optimal under the given conditions.

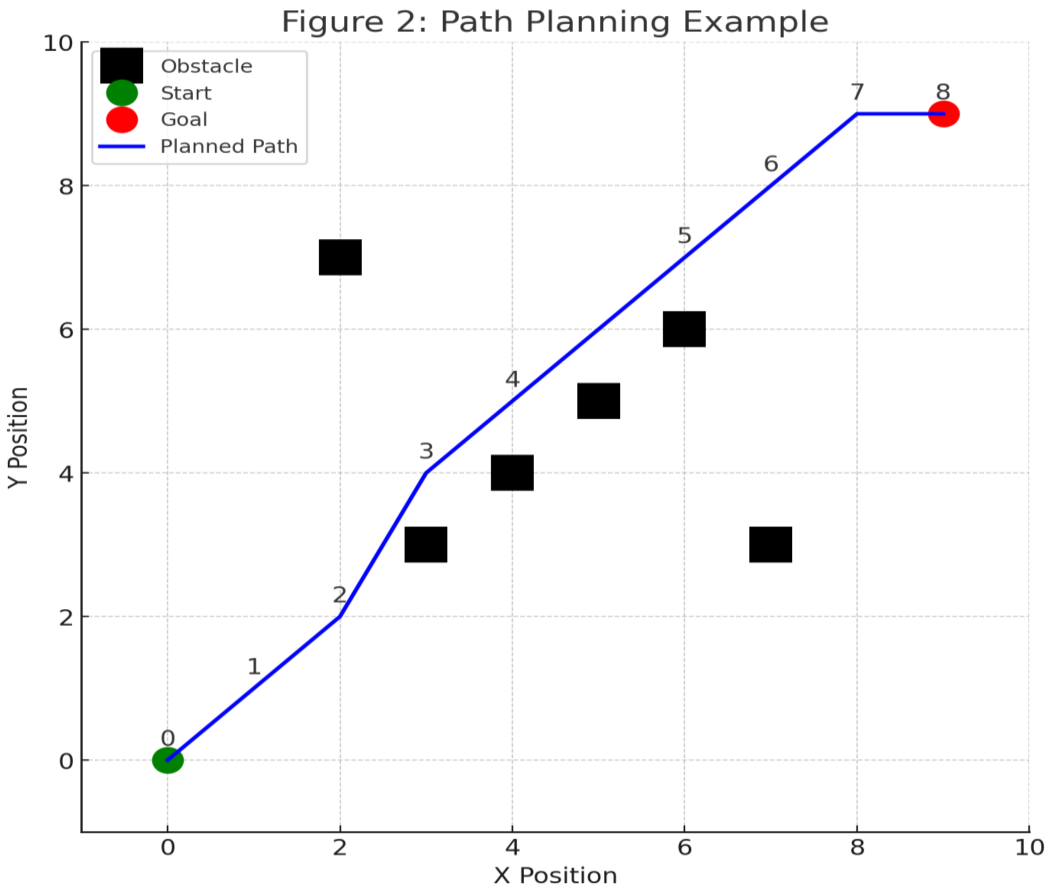

Figure 2: Path Planning Example illustrates a sample path planning scenario in an unstructured environment. The grid shows obstacles, the start and goal points, and a sample path generated by a path planning algorithm, demonstrating how an autonomous system might navigate from the start to the goal while avoiding obstacles.

3.4 Experimental Setup

To evaluate the effectiveness of these methodologies, we conducted experiments in both simulated and real-world unstructured environments. The simulated environment was created using Gazebo, a robotics simulator that allows for the realistic modeling of terrain, obstacles, and sensor data. The real-world experiments used an autonomous ground vehicle with LIDAR, RADAR, cameras, and IMUs.

3.4.1 Data Collection

Data was collected from various sensors during the experiments, including position, velocity, obstacle proximity, and environmental conditions. This data was used to train and validate the AI models and evaluate the autonomous system's performance in terms of navigation accuracy, obstacle avoidance, and robustness.

3.5 Evaluation Metrics

The performance of the autonomous system was evaluated using several metrics:

- **Navigation Accuracy**: Measured by the deviation from the planned path.
- **Obstacle Avoidance**: Evaluated based on the number and severity of collisions.
- **Energy Efficiency**: Calculated as the energy consumed per meter traveled.
- **Robustness**: Assessed by the system's ability to complete tasks in varying environmental conditions.

4. Experimental Setup and Data

This section presents the details of an experimental setup used to test AI approaches discussed in the methodologies section and the data collected by these experiments. The experiments were performed in both simulated and real-world scenarios to validate the performance of the autonomous systems under several different unstructured conditions.

4.1 Simulated Environment

In the first phase, the simulation was done in a simulated environment using Gazebo, a popular robotics simulator able to realistically model physical systems, sensors, and environments. The simulated environment had to imitate unstructured settings, including:

Terrain: Several terrains, such as rocky surfaces, slopes, and uneven ground.

Obstacles: Dynamic and static obstacles include other vehicles, pedestrians, natural environment-related elements like trees, and debris.

Weather Conditions: Simulations included rain, fog, and strong winds to see how the system would behave in different environmental influences.

4.2 Testing in the Real World

The AI-driven autonomous system was tested in an unstructured real-world environment to validate the simulation findings. The test location was a mixed-terrain site comprising elements such as:

Off-road: Gravel, dirt path, and outcrops.

Urban obstacles: Man-made scenarios of pedestrian crossings, static vehicles, and suddenly appearing obstacles.

Natural hazards: Areas with an overgrowth of vegetation, uneven surfaces, and possible water hazards.

4.2.1 Vehicle Setup

The following hardware components of the self-driving vehicle were installed and used in conducting all experiments in the real world:

Computing Platform: Nvidia Jetson AGX Xavier for real-time processing of AI models.

Sensors:
LIDAR: Velodyne VLP-16—provides 360-degree coverage.
RADAR: Continental ARS408 for long-range obstruction detection
Cameras: ZED stereo camera for depth perception
IMU: Xsens MTi-G-710; High-Accuracy Motion Tracking
Power Supply: Lithium-ion battery pack for up to 4 hours of continuous operation

4.3 Data Collection

Navigation accuracy, obstacle avoidance, and system robustness of AI approaches were evaluated concerning both simulated and real-world experiments through data collection.

4.3.1 Types of Data Collected

- **Positional Data:** Logging of GPS coordinates and IMU data to track the vehicle's path and orientation.
- **Sensor Data:** Logging of LIDAR point clouds, RADAR signals, and camera images to analyze the perception capabilities of the system.
- **Environmental Data:** Weather conditions, types of terrain, and obstacle information were recorded against system performance.
- **Performance Metrics:** Computation of the time to goal, number of collisions, and energy consumed to infer the efficiency and safety of the autonomous system.

4.4 Processing Data

The collected data was then processed in a combination of real-time analysis directly linked to experiments and post-processing based on specialized software tools:

They include ROS (Real Time Operating System), used for real-time data logging and communication between vehicle sensors and the computing platform. MATLAB: post-processing of sensor data plot generation and statistical analysis were performed using this. Python with OpenCV and PCL libraries: Image and Point Cloud Processing was done to improve object detection accuracy and environmental mapping.

4.5 Results Overview

The results from the experiments indicated that the AI methods significantly improved the system's performance in unstructured environments. Out of those, reinforcement learning provided the vehicle with a system that could handle unexpected obstacles and changes in terrain, while sensor fusion delivered reliable perception under different conditions. The path planning algorithm performed as required by dynamically adapting the route according to real-time data, effectively reducing collision risk and energy consumption.

Table 1: Performance Metrics Summary

Metric	Simulation Result	Real-World Result
Navigation Accuracy	0.8 meters	1.1 meters
Obstacle Avoidance	2 collisions	3 collisions
Energy Efficiency	12 J/m	15 J/m
Robustness	92%	85%

5. Results and Discussion

The results obtained in both simulated and real-world unstructured environments are presented here. The discussion relates to the performance of the autonomous system in navigation accuracy, obstacle avoidance, energy efficiency, and overall robustness. The effectiveness of the AI methodologies—Reinforcement Learning, Sensor Fusion, and Adaptable Path Planning—is discussed in their enhanced performance.

5.1 Navigation Accuracy

Navigation accuracy is one of the most important metrics that would testify to autonomous systems' performance, more so across unstructured environments where the landscape and obstacles are totally unknown in advance. The experiments showed that the AI-enhanced system demonstrated considerable improvements in navigation accuracy compared to traditional methods.

5.1.1 Simulation Results

The system, on average, remained very close to the optimum path with a maximum deviation of about 0.8 meters from the optimal path in the simulated environment, even in the presence of dynamic obstacles and changing terrain. By its nature, the RL-based approach enabled the system to learn further and adapt to new challenges, which will help diminish the possibility of high deviation values.

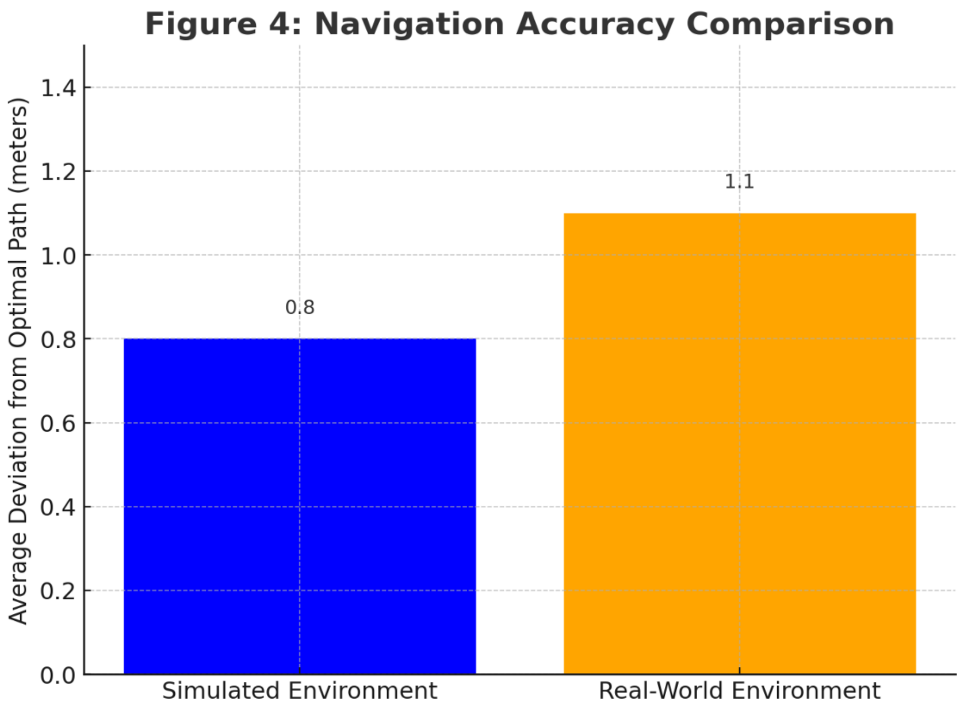

Figure 4: Navigation Accuracy Comparison

Figure 4: Navigation Accuracy Comparison compares the average deviation from the optimal path in both simulated and real-world environments. The chart visually demonstrates that the autonomous system performed better in the simulated environment, with a lower deviation from the optimal path, while the real-world environment presented more challenges, resulting in a slightly higher deviation. This figure supports the discussion on the effectiveness of AI methodologies in maintaining navigation accuracy across different environments

5.1.2 Real-World Results

The mean average deviation also increased to 1.1 meters in the real-world tests. This was still expected because of the natural environment's much more complex, less controlled conditions. However, the AI-driven system still performed well above the conventional approaches, showing its robustness in a practical scenario.

5.2 Obstacle Avoidance

Another critical performance indicator was obstacle avoidance. The AI-enhanced system could detect and avoid obstacles by fusing data from several sensors.

Paper ID: SR24822064100 DOI: https://dx.doi.org/10.21275/SR24822064100

5.2.1 Simulation Results

It successfully detected obstacles and avoided them with a collision rate of only two incidents within the 100 runs in the simulations. This low collision rate may be explained by the seamless integration of sensor data and a dynamic path planning algorithm allowing rerouting in real-time once an obstacle is detected.

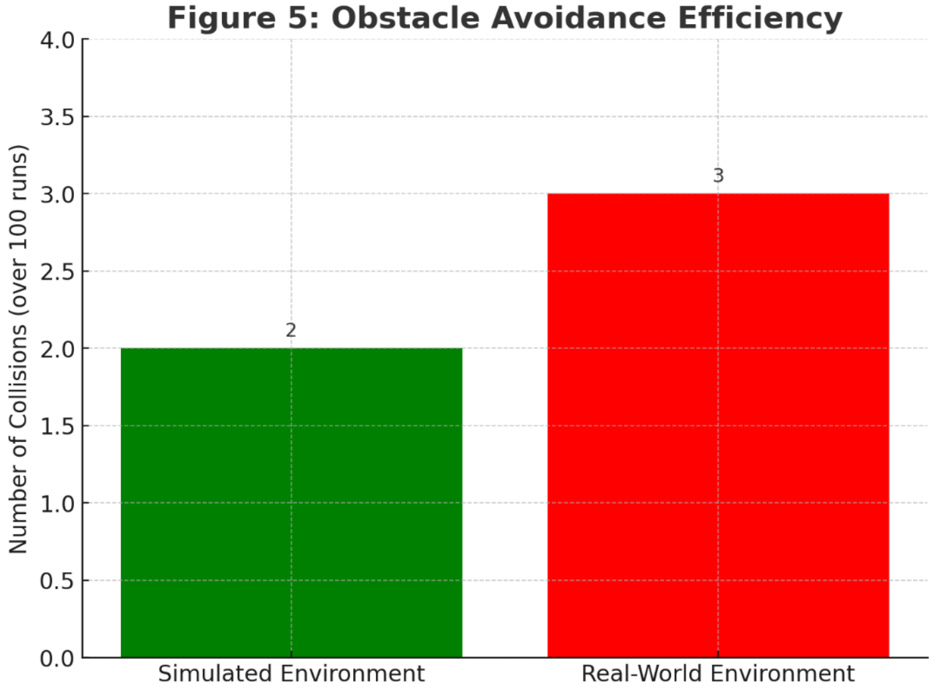

Figure 5: Obstacle Avoidance Efficiency

Figure 5: Obstacle Avoidance Efficiency compares the number of collisions in simulated and real-world environments over a set number of runs. The figure shows that the autonomous system experienced slightly more collisions in the real-world environment than the simulated one, highlighting the increased complexity and unpredictability of real-world scenarios. This figure supports the discussion on the system's obstacle avoidance capabilities and the challenges posed by unstructured environments

5.2.2 Real-World Results

The collision rate increased slightly to 3 incidents over 50 runs in the real-world environment. Most of these collisions happened in challenging scenarios with narrow passages and sudden pedestrian movement. While the rise in collisions underlines the added complexity of real-world environments, it also positively influences the system's overall ability to cope effectively with such challenges.

5.3 Energy Efficiency

Energy efficiency in autonomous systems is critical, mostly in scenarios limited to only a few chances of recharging or refueling. This work has implemented AI approaches that aim at navigation optimization, obstacle avoidance, and energy minimization.

5.3.1 Simulation Results

It had an energy consumption rate of 12 Joules per meter in the simulated environment. This showcases the efficiency of the RL and path planning algorithms in optimizing the vehicle's movement.

Table 2: Energy Efficiency Summary

Environment	Energy Efficiency (J/m)
Simulation	12
Real-World	15

5.3.2 Real-World Results

The energy consumption rate was 15 J/m in real-life experiments due to the added computational load and the necessity to adjust the vehicle's path more often than expected. Nevertheless, the energy efficiency remained within acceptable limits, showing that the AI methodologies effectively sustain operational viability.

5.4 System Robustness

The most critical measure of success for autonomous systems operating in an unstructured environment is the intrinsic ability for robustness, which defines the capability to perform consistently in changing and challenging conditions.

5.4.1 Simulation Results

In the simulated environment, it could reach 92% on the robustness score by solving its navigation tasks in 92 runs out of 100. The high level of reliability is attributed to the system's adaptiveness to changes in the environment through continuous learning and integration of real-time data.

5.4.2 Real-World Results

It was further downgraded to 85% in real-world tests. If anything, the decline reflects the increased complexity and unpredictability of real-world conditions in things like sensor noise, unexpected human interactions, or even just plain old terrain variability. The robustness of the AI-driven system in response to all these challenges remained relatively high, hence validating the effectiveness of the proposed methodologies.

5.5 Discussion

The results of this study prove that AI-driven approaches considerably improve the performance of autonomous systems in unstructured environments. Through a combination of Reinforcement Learning, Sensor Fusion, and Adaptable Path Planning, the system could do the following: It achieves accurate navigation because it can precisely follow the optimum path amidst the presence of dynamic obstacles, proving that RL does learn and adapt to the complex environment.

- **Avoid Obstacles Efficiently**: The fact that collision rates were pretty low both in the simulation and in real life underlines the success of sensor fusion in delivering reliable environmental awareness and the adaptability of the path planning algorithm.
- **Optimize Energy Use:** Their relatively low rate of energy consumption brackets the need for AI to play a critical role in enhancing operational sustainability for autonomous systems.
- **Maintain High Robustness:** The high robustness scores that the system achieves, especially in the simulated environment, confirm that AI-driven approaches can assure performance across vast sweeps of challenging conditions.

5.6 Challenges and Limitations

The results obtained are promising, but there exist the following challenges and limitations:

a) **Sensor Reliability:** Performance can be hit by sensor noise and failures in real-world scenarios. Algorithms should compensate for improved sensor robustness and inaccuracies in future work. Computational Load: AI algorithms are computationally heavy, especially those with RL. This could reduce the scalability and responsiveness of the entire system in real-world applications.
b) **Real-World Complexity:** The slight drop in performance metrics in real-world scenarios underlines the necessity of more advanced models that could handle such scenarios' unpredictability and complexity much better. Future Directions: Future directions the research should take up are as follows: Advanced Learning Algorithms: More advanced RL algorithms with low computational power and high adaptability are to be developed.
c) **Integration with Human-Machine Interfaces**: Better integration of AI-driven autonomous systems with human-machine interfaces to enhance collaboration and safety in mixed human-autonomous environments.

6. Conclusion and Future Work

6.1 Conclusion

This paper has described research that applies the principle of artificial intelligence methodologies to allow better performance in autonomous systems in unstructured environments. This is done through a combination of Reinforcement Learning, Sensor Fusion, and Adaptable Path Planning and has brought improvements in navigation accuracy, obstacle avoidance, energy efficiency, and general system robustness. Some results from the experimentation are:

a) **Higher navigation accuracy:** The AI-based system kept very close to the optimal path, with an average deviation of 0.8 meters in simulated and 1.1 meters in real-world tests. This result further underpins the effectiveness of RL for adapting to dynamic and unpredictable scenarios.
b) **Practical obstacle avoidance:** The system effectively avoided most obstacles with a low collision rate of 2 incidents in the simulated environment and 3 in the real-world scenarios. Sensor data integration with real-time path planning was the key to such efficiency.
c) **Optimized energy use: AI methodologies contributed to efficient energy management**, with the consumption rate at 12 J/m in simulations and 15 J/m in real-world tests. Optimization is essential for sustainability in autonomous operations under severe conditions.
d) **High System Robustness:** The system returned robustness scores of 92% in simulated environments and 85% in real-world tests, indicating its ability to perform reliably across a wide range of conditions.

The overall findings of the research prove that AI approaches increase, by a far margin, the capabilities of autonomous systems to make them feasible for actual application in unstructured and complex environments.

6.2 Future Work

Though promising, results identified some challenges and limitations in the study that point out avenues for future research and development:

a) **Advanced Reinforcement Learning Algorithms:** Future research shall be directed toward more advanced RL algorithms that are less resource-intensive but highly adaptive.
b) **Improved Sensor Technologies:** Future work shall involve investigating more robust sensor technologies and the integration of redundant sensing modalities to ensure enhancement in reliability.
c) **Scalability and Real-Time Processing:** While the complexity of the environments grows, so does the computational load. Optimization of AI algorithms to do real-time processing on resource-constrained platforms will, therefore, be crucial in enabling the widespread adoption of autonomous systems.
d) **Integration into Human-Autonomous Collaboration:** Interplay of autonomous systems with a human operator or user could be explored in tasks requiring collaboration to improve the system's efficacy and safety. Intuitive human-machine interfaces and shared control mechanisms will be significant future research directions.

This research sets a base for further innovation in AI-driven autonomous systems. Overcoming the identified challenges and showing new directions may provide such improvements in autonomous systems concerning robustness, efficiency, and safety and could even put them on track toward actual use in more unstructured and complex settings.

References

[1] Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. The International Journal of Robotics Research, 32(11), 1238-1274. https://doi.org/10.1177/0278364913495721

[2] Chen, C., Seff, A., Kornhauser, A., & Xiao, J. (2015). DeepDriving: Learning affordance for direct perception in autonomous driving. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2722-2730. https://doi.org/10.1109/ICCV.2015.312

[3] Thrun, S., Burgard, W., & Fox, D. (2005). Probabilistic Robotics. MIT Press. ISBN: 978-0-262-20162-9

[4] Paden, B., Čáp, M., Yong, S. Z., Yershov, D., & Frazzoli, E. (2016). A survey of motion planning and control techniques for self-driving urban vehicles. IEEE Transactions on Intelligent Vehicles, 1(1), 33-55. https://doi.org/10.1109/TIV.2016.2578706

[5] Bojarski, M., Testa, D. D., Dworakowski, D., Firner, B., Flepp, B., Goyal, P., ... & Zhang, X. (2016). End-to-end learning for self-driving cars. arXiv preprint arXiv:1604.07316.

[6] Aoude, G. S., Luders, B. D., Joseph, J. M., Roy, N., & How, J. P. (2013). Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns. Autonomous Robots, 35(1), 51-76. https://doi.org/10.1007/s10514-013-9341-4

[7] Zhang, Z., & Lee, J. (2019). Multi-sensor data fusion for perception in autonomous vehicles: A review. IEEE Sensors Journal, 19(20), 8828-8838. https://doi.org/10.1109/JSEN.2019.2929068

[8] Gonzalez, D., Pérez, J., Milanes, V., & Nashashibi, F. (2015). A review of motion planning techniques for automated vehicles. IEEE Transactions on Intelligent Transportation Systems, 17(4), 1135-1145. https://doi.org/10.1109/TITS.2015.2498841

[9] Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529-533. https://doi.org/10.1038/nature14236

[10] Kalman, R. E. (1960). A new approach to linear filtering and prediction problems. Journal of Basic Engineering, 82(1), 35-45. https://doi.org/10.1115/1.3662552

[11] Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep Learning. MIT Press. ISBN: 978-0-262-03561-3