

Performance Comparison of Duel-DDPG and DDPG Algorithms in the Decision-making Phase of Autonomous Vehicles

Ali Rizehvandi^{1*}, Shahram Azadi¹, Arno Eichberger²

¹ Faculty of Mechanical Engineering, K. N. Toosi University of Technology, 7, Pardis Ave., Mollasadra Str., 15418-49611 Tehran, Iran

² Institute of Automotive Engineering, Graz University of Technology, Inffeldgasse 11/2, 8010 Graz, Austria

* Corresponding author, e-mail: Ali.Rizehvandi@email.kntu.ac.ir

Received: 03 January 2025, Accepted: 26 October 2025, Published online: 15 December 2025

Abstract

Automated driving (AD) is a developing technology aimed at decreasing traffic accidents and enhancing driving efficiency. This research seeks to create a decision-making approach for self-driving cars, emphasizing actions such as changing lanes, overtaking, and maintaining lane position on highways, through deep reinforcement learning (DRL). In order to achieve this, a driving environment simulating a highway is established in the commercial multi-body simulation software IPG CarMaker 11, allowing the ego vehicle to navigate around other vehicles safely and efficiently. A control framework with a hierarchical structure is established to oversee these vehicles, where the high-level control is tasked with making driving choices. Additionally, the Duel Deep Deterministic Policy Gradient (Duel-DDPG) algorithm, which is a Deep Reinforcement Learning (DRL) method, is employed to create the highway decision-making strategy, which is simulated using MATLAB software. The computational methods of the Duel-DDPG and DDPG algorithms are examined and contrasted. A series of simulation evaluations are performed to evaluate the efficacy of the suggested decision-making policy. The results emphasize the advantages of the proposed framework regarding convergence rate and control effectiveness. The findings indicate that the Duel-DDPG-based approach effectively and safely performs highway driving activities.

Keywords

duel deep deterministic policy gradient algorithm, decision-making, deep reinforcement learning, IPG CarMaker software

1 Introduction

Autonomous vehicles (AVs) are revolutionizing transportation by allowing vehicles to navigate a wide variety of driving conditions without human intervention, leveraging the vast potential of artificial intelligence (AI) (Raj et al., 2020, Liu et al., 2019, Rasouli and Tsotsos, 2020). Companies like Toyota, Tesla, Ford, and Waymo are at the forefront of these advancements, focusing on key components essential for AV success: perception, decision-making, planning, and control (Gkartzonikas and Gkritza 2019). Perception involves the use of sensors such as lidar, radar, cameras, and GPS to detect the vehicle's environment. The decision-making system governs driving behaviors like lane changes, acceleration, and braking, while the planning module helps determine the optimal route. Control ensures that the vehicle executes these maneuvers accurately (Hoel, et al., 2020).

Autonomous vehicles are classified into six different levels, from Level 0 to Level 5, based on the complexity of their capabilities.

A crucial element of AV development is decision-making, which is often modeled on human cognitive functions. These decision-making frameworks can be derived from human driving experiences or through advanced learning methods. Reinforcement learning (RL), especially deep reinforcement learning (DRL), has demonstrated considerable potential in solving decision-making problems in autonomous driving (AD) (Sakib, 2020). For example, in Alizadeh et al. (2019), deep Q-learning (DQL) was employed to handle lane-changing choices in unpredictable highway situations. In a similar vein, Zhang et al. (2019) proposed an exploration policy based on a model for lane-changing, which included intrinsic rewards for unexpected outcomes. They offered a comprehensive summary of reinforcement learning (RL) and deep reinforcement learning (DRL) implementations in self-driving cars, discussing agent training, assessment methods, and reliable estimation (Kiran et al., 2020). Nonetheless, DRL-driven decision-making approaches encounter various obstacles, including problems related to sample efficiency,

prolonged learning periods, and safety during operation, which hinder their practical use in real scenarios.

In Furda and Vlacic (2011), a sophisticated decision-making system was created for urban traffic situations, encompassing various criteria to assist city vehicles in making effective decisions across different scenarios. Furthermore, Nie et al. (2016) investigated a strategy for lane-changing decisions in connected autonomous vehicles, combining cooperative car-following models with a candidate decision-making module. Additionally, in Li et al. (2018), the researchers presented a driving system that mimics human behavior by modifying its driving choices based on the preferences of individual drivers.

Deep reinforcement learning (DRL) methods are being increasingly utilized to address intricate sequential decision-making challenges in autonomous driving (AD). For instance, Duan et al. (2020) created a hierarchical framework for acquiring decision-making policies through reinforcement learning (RL), removing the necessity for previously labeled driving data. DRL has emerged as a valuable resource for addressing long-term sequential decision-making challenges. Lately, a variety of research has investigated DRL in the context of automated driving. For example, Duane and his team introduced a hierarchical framework for developing decision-making policies by utilizing RL methods. Other studies (Yang et al., 2020a, Mnih et al., 2015) have applied DRL to tackle difficulties such as preventing collisions and arranging trails in autonomous vehicles, showing that DRL outperforms traditional RL methods. Furthermore, studies (Chen et al., 2020, Yang, et al., 2020b) expanded the scope to include fuel consumption optimization, by utilizing the deep Q-learning (DQL) algorithm, which has shown to be effective for driving-related tasks. Han and Miao (2020) implemented DQL for making decisions related to lane changing and lane keeping in connected autonomous vehicles, using data from surrounding vehicles as input. Nevertheless, traditional DRL approaches struggle with highway overtaking tasks owing to the continuous action space and extensive state space.

In Nagesh Rao et al. (2019), the application of RL in unmanned driving has gained popularity; however, enhancing the stability of autonomous vehicles while maintaining effective path tracking and obstacle avoidance in diverse conditions continues to pose difficulties. To address the functional needs of path tracking and obstacle avoidance, a control strategy based on the deep deterministic policy gradient (DDPG) was introduced for unmanned vehicles to prevent collisions.

In Lv et al. (2022), a motion planning approach based on deep reinforcement learning (DRL) was proposed for autonomous driving applications in highway environments, focusing on situations where an autonomous vehicle merges into a two-lane traffic flow and executes lane-change actions. The research combined a DRL model with the autonomous driving system through an end-to-end learning framework, utilizing the deep deterministic policy gradient (DDPG) algorithm along with defined reward functions.

The objective of this research is to create a decision-making policy that guarantees both effectiveness and safety for autonomous vehicles on highways. To accomplish this, the research introduces a DRL approach enhanced by Duel-DDPG. This is the first time the Duel-DDPG algorithm is not only applied but also used to solve highway navigation problems involving long driving scenarios, including lane changes, double lane changes, and lane keeping. One significant benefit of employing the Duel-DDPG algorithm in autonomous driving is its effectiveness in functioning within continuous action spaces over long durations, as well as its ability to mimic human decision-making in prolonged situations efficiently. In contrast to conventional discrete action algorithms like deep Q-networks (DQN), which limit the agent to a predetermined set of actions, Duel-DDPG allows for more precise manipulation of the vehicle's steering, acceleration, and braking (AVL List GmbH., online). This level of precision is vital for navigating highways, where maintaining smooth and accurate control is important for both safety and comfort. Utilizing the actor-critic framework, Duel-DDPG can develop a strategy that produces continuous actions within intricate highway settings, facilitating more intuitive and effective driving choices similar to those made by humans. This leads to enhanced performance in dynamic and challenging driving scenarios, where distinct actions can result in less smooth and sharper movements.

As illustrated in Fig. 1, the process of decision-making on highways is conducted using the advanced DRL approach, particularly the Duel-DDPG algorithm. To facilitate this, the driving environment will be replicated through the CarMaker simulator, allowing the Duel-DDPG algorithm to acquire driving maneuvers such as steering, accelerating, and decelerating.

The research starts by outlining the actual vehicle dynamics and the driving scenarios within the CarMaker simulator (Reichmann-Blaga, 2024) to guarantee that the autonomous vehicle functions both safely and effectively. The Duel-DDPG enhanced DRL technique utilizes the actor-critic framework to directly acquire control actions and mimic

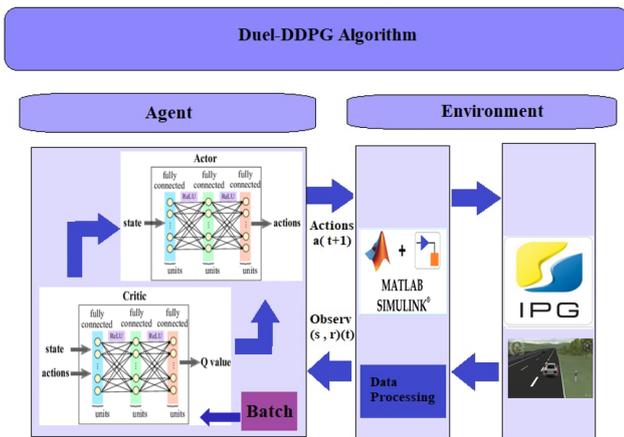


Fig. 1 A decision-making policy for autonomous vehicles, utilizing DRL method

human behavior, creating a trust region with adjusted objectives. The subsequent sections provide details on the implementation of the DRL algorithm. Finally, the effectiveness of the decision-making algorithm is assessed for particular scenarios, such as lane changes, double lane changes, and lane keeping, while also confirming the algorithm's flexibility in various situations. Actually, this study hypothesizes that the Duel-Deep Deterministic Policy Gradient (Duel-DDPG) algorithm will outperform the standard Deep Deterministic Policy Gradient (DDPG) algorithm in autonomous vehicle decision-making tasks, specifically demonstrating improved safety and effectiveness in highway driving scenarios through more efficient learning and better policy evaluation.

This paper introduces three key novelty aimed at improving the safety and efficiency of autonomous driving on highways in real-world scenarios. These contributions include:

1. The using of novel DRL method, namely the Duel-DDPG algorithm.
2. The development of an advanced, safe, and efficient decision-making policy for autonomous driving on highways using the Duel-DDPG algorithm.
3. The use of actual driving scenarios in the CarMaker simulator.

The current study is structured as follows to elaborate on these contributions:

- Section 2 specifies the research on Duel-DDPG enhanced DRL. Section 3 presents the evaluation of simulation results related to the proposed decision-making strategy. Finally, Section 4 provides the concluding remarks. This research utilized the vehicle dynamics model available through the commercial multi-body simulation tool IPG CarMaker.

- The subsequent section introduces a deep reinforcement learning (DRL) approach to facilitate the learning process and establish the highway decision policy.

2 Methodology

Machine learning (ML), which is a branch of artificial intelligence (AI), aims to enhance the effectiveness of computational algorithms through the use of data. Machine learning (ML) is typically categorized into three primary types: reinforcement learning (RL), supervised learning, and unsupervised learning. In the reinforcement learning framework, an independent agent acquires skills to carry out a task in a specific environment by striving to optimize a defined reward function. The agent receives rewards for making choices during its interactions with the environment, whereas making some choices leads to penalties or negative rewards.

Supervised learning consists of acquiring knowledge from examples that are labeled by experts. However, this approach is not ideal for addressing interactive problems, since obtaining labels that truly represent every potential interaction is challenging. In contrast, unsupervised learning aims to identify underlying patterns within unlabeled data. Although this method can reveal valuable structures in the data, it does not focus on maximizing a reward, which is the main objective of the RL.

In certain reinforcement learning scenarios, the environment might have a large number of states and actions (Hugging Face., online). In these situations, an artificial neural network (ANN) can serve as a function approximator. When an ANN is employed in reinforcement learning to estimate functions, it is known as DRL.

2.1 Duel Deep Deterministic Policy Gradient (Duel-DDPG) algorithm

One of the main challenges in Deep Deterministic Policy Gradient (DDPG) algorithms is the accuracy and efficiency of the critic network. This network needs to be able to accurately predict the value of each action in different situations. Improving the accuracy of the critic network can directly lead to better overall performance of the algorithm (Song et al. 2016).

The advantage network has been introduced as an improved method in reinforcement learning algorithms. The core concept of this network involves decomposing the overall value of a state into two components: the base value of the state and the advantage of an action. This decomposition can enhance the learning process and improve the accuracy of the critic network.

In this work, we use the advantage network within the structure of the critic network in the Deep Deterministic Policy Gradient algorithm. This innovation is designed to improve the accuracy and efficiency of the critic network, ultimately leading to enhanced overall performance of the DDPG algorithm in complex environments.

In this approach, instead of directly predicting the total value $Q(s, a)$, the critic network first predicts the base value of the state $V(s)$ and the advantage of the action $A(s, a)$ (as shown in Fig. 2). Then, the total value is calculated as the combination of these two values:

$$Q(s, a) = V(s) + A(s, a) \quad (1)$$

This structure can help reduce the variance in value predictions, leading to more stable and faster learning.

The Duel Deep Deterministic Policy Gradient (Duel-DDPG) algorithm is a groundbreaking reinforcement learning approach that combines both value-oriented and policy-oriented strategies. It is particularly adept at addressing challenges involving continuous action spaces in reinforcement learning and behaves as human in decision-making.

Also, the Duel-DDPG employs an actor-critic architecture, consisting of two key components:

- Actor-network: This network is designed to learn the policy function that relates states to actions. Its objective is to optimize the expected return by choosing actions based on the present state while also assessing the value of each action.
- Critic network: This network is responsible for learning the value function, which estimates the expected return (or total reward) associated with adhering to a specific policy. It aids in assessing the actions selected by the actor-network.

Furthermore, Duel-DDPG operates as an off-policy algorithm, which implies that it learns from data drawn from an experience replay buffer instead of adhering to a specific policy at all times. It is also categorized as model-free, signifying that it does not need to understand the fundamental dynamics of the environment. In contrast to algorithms that are optimal for discrete action spaces, Duel-DDPG is tailored for continuous action spaces, making it applicable to various fields, such as autonomous vehicle operations and control tasks.

The mean-squared Bellman error (MSBE) is defined as:

$$L(\theta, \theta) = E_D \left[\left(Q_\theta(s, a) - (r + \gamma(1-d) \max_{a'} Q_\theta(s, a')) \right)^2 \right] \quad (2)$$

The Duel-DDPG algorithm integrates elements from both policy gradient techniques and Q-learning. The actor network utilizes policy gradient methods to directly enhance the expected return, whereas the critic network employs Q-learning to estimate the value of state-action combinations.

Furthermore, Duel-DDPG utilizes target networks to enhance the stability of the training process. These target networks serve as duplicates of the actor and critic networks but are refreshed less often than the primary networks. Duel-DDPG also employs an experience replay buffer that retains and samples experiences throughout training. This practice aids in diversifying the data and improving sample efficiency.

Moreover, the reward function is structured as follows:

$$R_t = \alpha \cdot |\theta| + \beta \cdot |a| + \gamma \cdot \text{collision} + M \quad (3)$$

In this context, $\alpha=0.5$, $\beta=0.1$ and $\gamma=-1$ also, θ represents the steering angle, a denotes longitudinal acceleration, and a collision occurs when d_{relative} equals 0, with M being a constant value. According to Eq. (3), the agent is trained in a highway setting to maneuver across the highway exit.

Fig. 3 shows the designed decision-making algorithm in the Matlab - Simulink software.

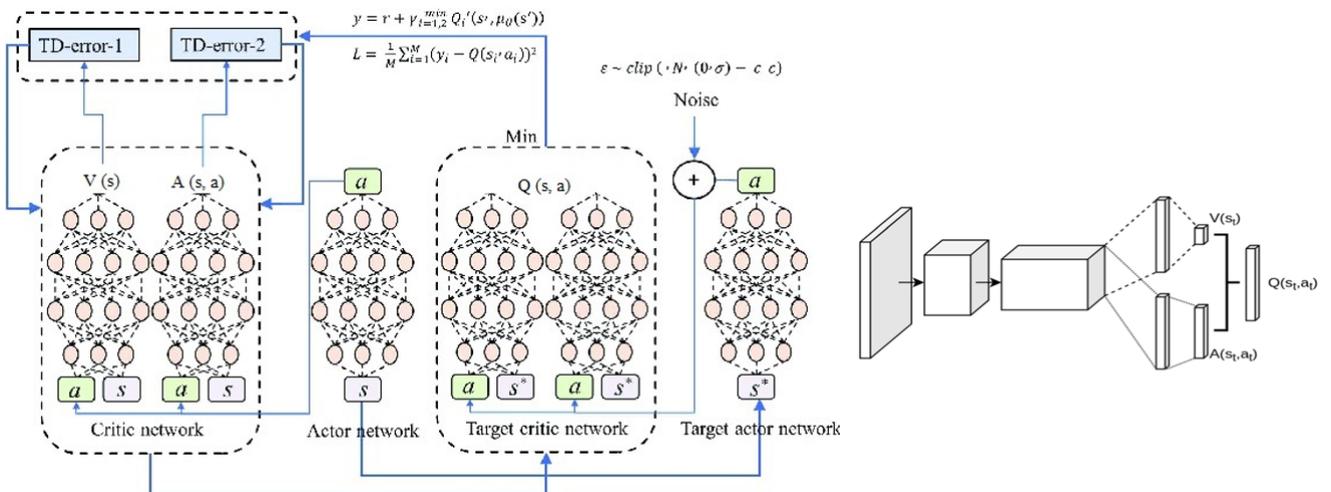


Fig. 2 The Duel Deep Deterministic Policy Gradient (Duel-DDPG) algorithm structure

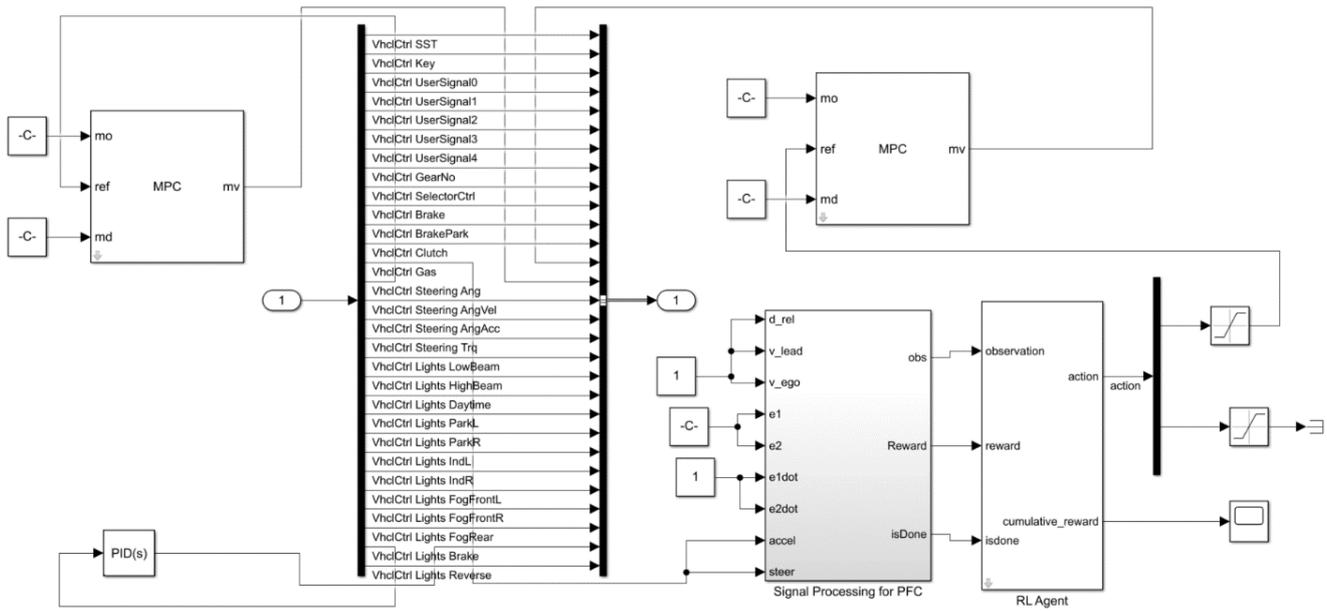


Fig. 3 The simulated Duel-DDPG algorithm in Simulink software (MathWorks Inc., online)

In the following section, we will evaluate and confirm the effectiveness and validity of the proposed decision-making algorithm.

3 Results and discussion

This section assesses the efficacy of the suggested decision-making strategy for the autonomous electric vehicle (AEV) utilizing the Duel Deep Deterministic Policy Gradient (Duel-DDPG) method. The assessment concentrates on three primary elements. Initially, it assesses and confirms the effectiveness of this decision-making method compared to another approach through comprehensive simulation results that demonstrate its advantages. Additionally, it verifies the learning ability of the Duel-DDPG algorithm by examining the total accumulated rewards. Finally, it explores the flexibility of the generated decision-making strategy by evaluating it in two similar highway driving situations. Fig. 3 depicts the particular scenario carried out by the Duel-DDPG agent in the CarMaker simulator.

In this study, the autonomous vehicle performs maneuvers such as lane changes, overtaking, and lane keeping, similar to adaptive cruise control systems, within a highway environment with complex traffic, in order to exit the highway. The objective is to evaluate the ability of the Duel-DDPG algorithm to navigate around moving obstacles in heavy traffic and perform maneuvers safely and efficiently to exit the highway (as shown in Fig. 4).

As depicted in Fig. 4, the scenario is implemented in the CarMaker simulator. In this scenario, the vehicle is traveling at 70 km/h, and considering the complex traffic environment

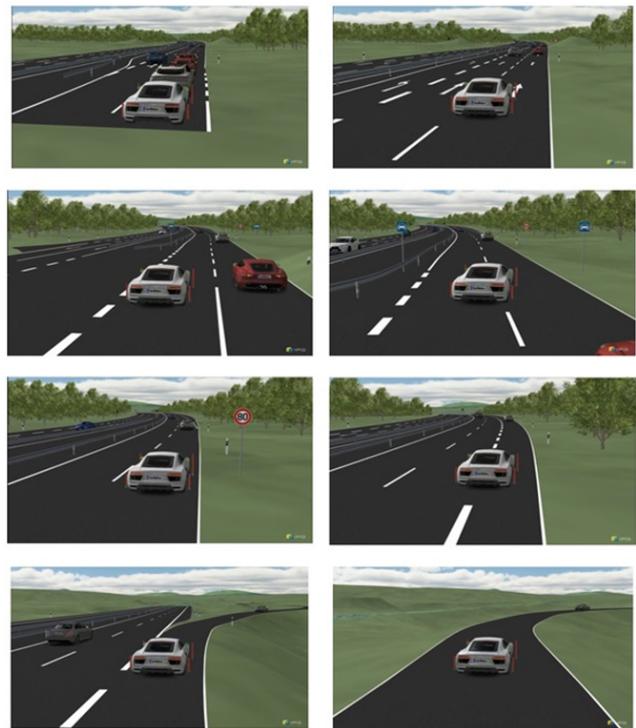


Fig. 4 Performing lane change, overtaking, and lane-keeping maneuvers similar to an adaptive cruise control system by the autonomous vehicle in the highway environment

it is navigating, it performs maneuvers such as lane changes, overtaking, and lane keeping like an adaptive cruise control system, safely passing through the traffic and approaching the highway exit before ultimately leaving it.

In this scenario, the Duel-DDPG algorithm successfully executed lane changes, overtaking, and lane-keeping maneuvers, similar to adaptive cruise control, without

colliding with other vehicles, while ensuring safety and efficiency, ultimately exiting the highway. Now, the analysis of the lateral acceleration, longitudinal speed, and yaw rate of the autonomous vehicle equipped with the Duel-DDPG and DDPG algorithms will be discussed.

As shown in Fig. 5, in the heavy traffic scenario, the average lateral acceleration in the Duel-DDPG algorithm was 0.5 m/s^2 , whereas in the DDPG algorithm, it was 1.1 m/s^2 . Additionally, the standard deviation of lateral acceleration in the Duel-DDPG algorithm was lower, indicating better vehicle control during maneuvers. Also, the reduction in lateral acceleration in the Duel-DDPG algorithm positively impacted the vehicle's stability and safety. This improvement, especially during lane changes and abrupt maneuvers, reduced unwanted deviations and provided better performance in lateral control. Moreover, the Duel critic network in the Duel-DDPG algorithm, through the decomposition of the Q-value function, was able to select better policies for controlling the vehicle's lateral movements. This optimization led to reduced lateral acceleration and improved vehicle control and precision.

The results indicate that in the simulated scenario, the Duel-DDPG algorithm, with more precise control and lower lateral acceleration, was able to guide the vehicle with smoother and more stable movements. This demonstrates the algorithm's higher capability when facing complex road conditions. In other words, the Duel-DDPG algorithm, by reducing lateral acceleration, was able to provide more precise and stable control in vehicle motion, enhancing safety while reducing unwanted lateral fluctuations, thus improving passenger comfort—particularly in the complex traffic conditions of the highway.

As seen in Fig. 6, the autonomous vehicle equipped with the Duel-DDPG algorithm, in response to the heavy traffic environment it faced, moved at a maximum set speed of 19.4 m/s (70 km/h) in the traffic-free environment. It performed more effectively and efficiently than the DDPG algorithm in the simulated scenario, successfully exiting the highway.

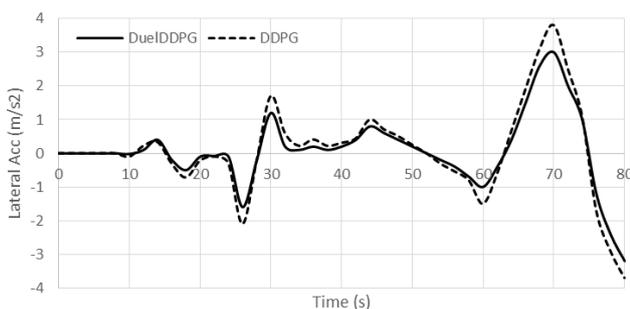


Fig. 5 Lateral acceleration in the Duel-DDPG and DDPG algorithms

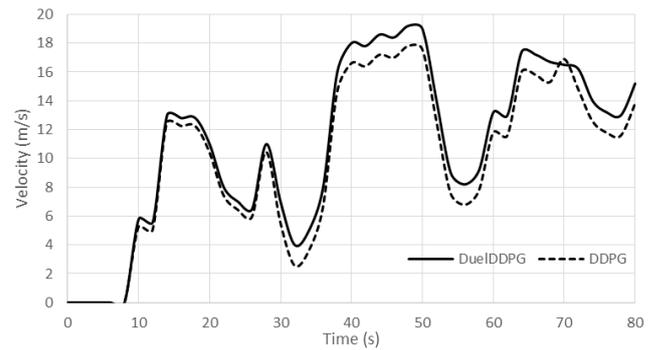


Fig. 6 Longitudinal speed in the Duel-DDPG and DDPG algorithms

As shown in Fig. 7, the Duel-DDPG algorithm effectively controlled the yaw rate in the modeled heavy traffic scenario, demonstrating its superiority over the DDPG algorithm in maintaining stability and executing precise maneuvers. This also indicates better decision-making alignment with real road conditions and vehicle movement. Simulation results also show that the average longitudinal acceleration in the Duel-DDPG algorithm was 1.2 m/s^2 , while in the DDPG algorithm, it was 0.9 m/s^2 . Additionally, the standard deviation of acceleration in the Duel-DDPG algorithm was lower, indicating greater system stability.

Actually, in the modeled heavy scenario, the Duel-DDPG algorithm responded faster than the DDPG algorithm and was able to provide more suitable longitudinal acceleration. This allowed the autonomous vehicle to reach the highway exit faster and exit safely. Furthermore, the Duel structure in the Duel-DDPG algorithm improved decision-making for longitudinal acceleration control. By decomposing the value function into base and advantage components, the algorithm was able to choose better policies, leading to increased acceleration and reduced time to reach optimal speeds in this heavy traffic scenario. The results show that even in a heavy traffic environment, the Duel-DDPG algorithm, due to its use of the Duel critic network, exhibited less fluctuation in longitudinal acceleration. This greater stability allows the vehicle to operate more safely in high-risk situations and critical decision-making scenarios. The improved

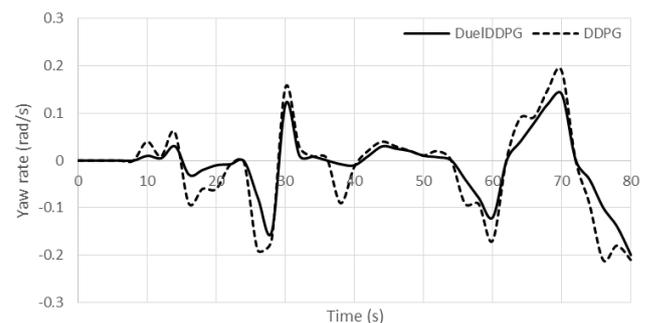


Fig. 7 Yaw rate in the Duel-DDPG and DDPG algorithms

longitudinal acceleration in the Duel-DDPG algorithm, compared to the DDPG algorithm, indicates that this algorithm not only performed better in complex decision-making but also in key operational parameters, such as longitudinal acceleration, even in heavy highway traffic conditions.

Based on the analysis, the autonomous vehicle equipped with the Duel-DDPG algorithm is able to effectively reduce its speed in heavy traffic and continue with lane-keeping maneuvers. When facing a clear lane, it adjusted its speed to the maximum setting of 70 km/h and safely exited the highway without any collisions. In this scenario, the Duel-DDPG algorithm successfully adjusted the vehicle's speed and continued the maneuver by performing lane changes, overtaking, and lane keeping. It also maintained the maximum speed of 70 km/h when encountering an unobstructed lane, ultimately completing the maneuver and exiting the highway.

4 Conclusion

The research utilizes Deep Reinforcement Learning (DRL) approaches to address the complexities associated with decision-making on highways. This investigation presents a novel strategy by integrating a distinctive set of driving scenarios, lane changing, overtaking, and lane keeping, which have not been previously examined. This integration facilitates a thorough assessment of the decision-making capabilities of the Duel-DDPG algorithm across diverse and intricate driving environments. By tackling this multifaceted scenario, the current study establishes a more practical testing framework, thereby improving the dependability and flexibility of autonomous driving (AD) systems in practical applications.

References

- Alizadeh, A., Moghadam, M., Bicer, Y., Ure, N., Yavas, U., Kurtulus, C. (2019) "Automated Lane Change Decision Making using Deep Reinforcement Learning in Dynamic and Uncertain Highway Environment", In: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), Auckland, New Zealand, pp. 1399–1404. ISBN 978-1-5386-7025-5
<https://doi.org/10.1109/ITSC.2019.8917192>
- AVL List GmbH (online) "Global Vehicle Benchmarking and Technology", [online] Available at: <https://www.avl.com/en/engineering/vehicle-engineering/vehicle-development/global-vehicle-benchmarkingand-technology> [Accessed: 17 November 2022]
- Chen, C., Jiang, J., Lv, N., Li, S. (2020) "An Intelligent Path Planning Scheme of Autonomous Vehicles Platoon Using Deep Reinforcement Learning on Network Edge", IEEE Access, 8, pp. 99059–99069.
<https://doi.org/10.1109/ACCESS.2020.2998015>
- Duan, J., Li, S. E., Guan, Y., Sun, Q., Cheng, B. (2020) "Hierarchical reinforcement learning for self-driving decision-making without reliance on labeled driving data", IET Intelligent Transport Systems, 14(5), pp. 297–305.
<https://doi.org/10.1049/iet-its.2019.0317>
- Furda, A., Vlacic, L. (2011) "Enabling Safe Autonomous Driving in Real-World City Traffic Using Multiple Criteria Decision Making", IEEE Intelligent Transportation Systems Magazine, 3(1), pp. 4–17.
<https://doi.org/10.1109/MITS.2011.940472>
- Gkartzonikas, C., Gkritza, K. (2019) "What have we learned? A review of stated preference and choice studies on autonomous vehicles", Transportation Research Part C: Emerging Technologies, 98, pp. 323–337.
<https://doi.org/10.1016/j.trc.2018.12.003>
- Han, S., Miao, F. (2020) "Behavior planning for connected autonomous vehicles using feedback deep reinforcement learning", [pdf] arXiv preprint. Available at: <http://arxiv.org/pdf/2003.04371v1> [Accessed: 4 September 2022].

This advancement not only addresses a significant gap in the current literature but also aids in the progression of safer and more effective AD technologies. Additionally, a tailored control framework is created using the Duel-DDPG algorithm in the given driving situations to ensure both safety and effectiveness. The document outlines the efficacy, speed of convergence, and flexibility of the suggested approach via a set of simulation experiments. The findings indicate that the Duel-DDPG algorithm demonstrates superior efficiency and safety compared to the DDPG technique. Additionally, the evaluation of the testing outcomes highlights the potential for the proposed method to be effectively applied in real-world driving situations. Upcoming research will concentrate on executing online decision-making for highways by utilizing hardware-in-the-loop experiments and harnessing actual highway databases to enhance pertinent overtaking techniques.

List of abbreviations

The following abbreviations are used in this manuscript:

| | |
|----------|---|
| AD | Automated Driving |
| DRL | Deep Reinforcement Learning |
| DuelDDPG | Duel Deep Deterministic Policy Gradient |
| DDPG | Deep Deterministic Policy Gradient |
| RL | Reinforcement Learning |
| DQL | Deep Q Learning |
| ML | Machine Learning |
| DQN | Deep Q-Network |

- Hoel, C.-J., Driggs-Campbell, K., Wolff, K., Laine, L., Kochenderfer, M. J. (2020) "Combining Planning and Deep Reinforcement Learning in Tactical Decision Making for Autonomous Driving", *IEEE Transactions on Intelligent Vehicles*, 5(2), pp. 294–305.
<https://doi.org/10.1109/TIV.2019.2955905>
- Hugging Face (online) "The Reinforcement Learning Framework—Hugging Face Deep RL Course", [online] Available at: <https://huggingface.co/learn/deep-rl-course/unit1/rl-framework> [Accessed: 4 May 2022].
- Kiran, B. R., Sobh, I., Talpaert, V., Mannion, P., Sallab, A. A., Yogamani, S., Pérez, P. (2020) "Deep Reinforcement Learning for Autonomous Driving: A Survey", [pdf] arXiv preprint.
<https://doi.org/10.48550/arXiv.2002.00444>
- Li, L., Ota, K., Dong, M. (2018) "Humanlike Driving: Empirical Decision-Making System for Autonomous Vehicles", *IEEE Transactions on Vehicular Technology*, 67(8), pp. 6814–6823.
<https://doi.org/10.1109/TVT.2018.2822762>
- Liu, T., Tian, B., Ai, Y., Chen, L., Liu, F., Cao, D. (2019) "Dynamic States Prediction in Autonomous Vehicles: Comparison of Three Different Methods", In: 2019 IEEE Intelligent Transportation Systems Conference (ITSC), Auckland, New Zealand, pp. 3750–3755. ISBN 978-1-5386-7025-5
<https://doi.org/10.1109/ITSC.2019.8916969>
- Lv, K., Pei, X., Chen, C., Xu, J. (2022) "A Safe and Efficient Lane Change Decision-Making Strategy of Autonomous Driving Based on Deep Reinforcement Learning", *Mathematics*, 10(9), 1551.
<https://doi.org/10.3390/math10091551>
- MathWorks Inc. (online) "MATLAB and Simulink software, (Version R2022a)", [computer program] Available at: <https://www.mathworks.com/products/simulink.html> [Accessed: 7 August 2025]
- 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, pp. 529–533.
<https://doi.org/10.1038/nature14236>
- Nagesh Rao, S., Tseng, H. E., Filev, D. (2019) "Autonomous Highway Driving using Deep Reinforcement Learning", In: 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC), Bari, Italy, pp. 2326–2331. ISBN 978-1-7281-4568-6
<https://doi.org/10.1109/SMC.2019.8914621>
- Nie, J., Zhang, J., Ding, W., Wan, X., Chen, X., Ran, B. (2016) "Decentralized Cooperative Lane-Changing Decision-Making for Connected Autonomous Vehicles", *IEEE Access*, 4, pp. 9413–9420.
<https://doi.org/10.1109/ACCESS.2017.2649567>
- Raj, A., Kumar, J. A., Bansal, P. (2020) "A multicriteria decision making approach to study barriers to the adoption of autonomous vehicles", *Transportation Research Part A: Policy and Practice*, 133, pp.122–137.
<https://doi.org/10.1016/j.tra.2020.01.013>
- Rasouli, A., Tsotsos, J. K. (2020) "Autonomous Vehicles That Interact With Pedestrians: A survey of Theory and Practice", *IEEE Transactions on Intelligent Transportation Systems*, 21(3), pp. 900–918.
<https://doi.org/10.1109/TITS.2019.2901817>
- Reichmann-Blaga, E. (2024) "Validierung von Fahrzeugdynamischen Simulationsmodellen anhand von 546 Fahrzeugmessungen" (Validation of vehicle dynamic simulation models based on 546 vehicle measurements), Master's Thesis, Graz University of Technology.
<https://doi.org/10.3217/3ktg-zmr02>
- Sakib, N. (2020) "Highway Lane change under uncertainty with Deep Reinforcement Learning based motion planner", Master's Thesis, University of Alberta.
<https://doi.org/10.7939/r3-qm5k-s682>
- Song, W., Xiong, G., Chen, H. (2016) "Intention-Aware Autonomous Driving Decision-Making in an Uncontrolled Intersection", *Mathematical Problems in Engineering*, 2016(1), 1025349.
<https://doi.org/10.1155/2016/1025349>
- Yang, C., Zha, M., Wang, W., Liu, K., Xiang, C. (2020a) "Efficient energy management strategy for hybrid electric vehicles/plug-in hybrid electric vehicles: review and recent advances under intelligent transportation system", *IET Intelligent Transport Systems*, 14(7), pp. 702–711.
<https://doi.org/10.1049/iet-its.2019.0606>
- Yang, C., You, S., Wang, W., Li, L., Xiang, C. (2020b) "A Stochastic Predictive Energy Management Strategy for Plug-in Hybrid Electric Vehicles Based on Fast Rolling Optimization", *IEEE Transactions on Industrial Electronics*, 67(11), pp. 9659–9670.
<https://doi.org/10.1109/TIE.2019.2955398>
- Zhang, S., Peng, H., Nagesh Rao, S., Tseng, E. (2019) "Discretionary Lane Change Decision Making using Reinforcement Learning with Model-Based Exploration", In: 2019 18th IEEE International Conference on Machine Learning and Applications (ICMLA), Boca Raton, FL, USA, pp. 844–850. ISBN 978-1-7281-4551-8
<https://doi.org/10.1109/ICMLA.2019.00147>