



HAL
open science

Behavioral decision-making for urban autonomous driving in the presence of pedestrians using Deep Recurrent Q-Network

Niranjan Deshpande, Dominique Vaufreydaz, Anne Spalanzani

► **To cite this version:**

Niranjan Deshpande, Dominique Vaufreydaz, Anne Spalanzani. Behavioral decision-making for urban autonomous driving in the presence of pedestrians using Deep Recurrent Q-Network. ICARCV - 16th International Conference on Control, Automation, Robotics and Vision, Dec 2020, Shenzhen, China. pp.1-9. hal-02977009

HAL Id: hal-02977009

<https://inria.hal.science/hal-02977009v1>

Submitted on 23 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Behavioral decision-making for urban autonomous driving in the presence of pedestrians using Deep Recurrent Q-Network

AUTHOR VERSION

Niranjan Deshpande

Univ. Grenoble Alpes, Inria, 38000 Grenoble, France
Niranjan.Deshpande@inria.fr

Dominique Vaufreydaz

Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, 38000 Grenoble, France
Dominique.Vaufreydaz@univ-grenoble-alpes.fr

Anne.Spalanzani@inria.fr

Univ. Grenoble Alpes, Inria, Grenoble INP, 38000 Grenoble, France
Anne.Spalanzani@inria.fr

ABSTRACT

Decision making for autonomous driving in urban environments is challenging due to the complexity of the road structure and the uncertainty in the behavior of diverse road users. Traditional methods consist of manually designed rules as the driving policy, which require expert domain knowledge, are difficult to generalize and might give sub-optimal results as the environment gets complex. Whereas, using reinforcement learning, optimal driving policy could be learned and improved automatically through several interactions with the environment. However, current research in the field of reinforcement learning for autonomous driving is mainly focused on highway setup with little to no emphasis on urban environments. In this work, a deep reinforcement learning based decision-making approach for high-level driving behavior is proposed for urban environments in the presence of pedestrians. For this, the use of Deep Recurrent Q-Network (DRQN) is explored, a method combining state-of-the art Deep Q-Network (DQN) with a long term short term memory (LSTM) layer helping the agent gain a memory of the environment. A 3-D state representation is designed as the input combined with a well defined reward function to train the agent for learning an appropriate behavior policy in a real-world like urban simulator. The proposed method is evaluated for dense urban scenarios and compared with a rule-based approach and results show that the proposed DRQN based driving behavior decision maker outperforms the rule-based approach.

Keywords Deep Recurrent Q-Network, Autonomous vehicle

1 INTRODUCTION

An autonomous driving system consists of two main sub-systems, namely: a perception sub-system and a navigation sub-system. The navigation sub-system is responsible for moving the autonomous vehicle from a source point A in the environment to a destination point B . For this, the navigation sub-system usually uses a decision making module, a planning and control module and an actuator mechanism. The task of decision making is to make an appropriate driving decision according to the information received from the perception sub-system about the surrounding environment and then a plan for drivable path is generated by a path planner which is further passed to the control module. Reasonable decision-making in a variety of complex environments is a great challenge, and it is impossible to enumerate coping strategies in various situations. Hence, a method that can learn a suitable behavior from its own experiences would be desirable.

Reinforcement learning (RL), a paradigm of machine learning, learns a policy automatically without expert domain knowledge. This is done by simulating various types of different scenarios, even some dangerous ones. Deep learning (DL) combined with reinforcement learning (RL) have led to the emergence of deep reinforcement learning (deep RL) based approaches, with Deep Q-Networks (DQN) being one of most commonly used deep RL algorithm. It has demonstrated good results in many areas such as video games [1], control theory [2] and robotics [3]. As a result, many deep RL based methods are being presented for autonomous driving as well. However, most of these approaches focus on highway driving scenarios and with limited to no focus on complex urban scenarios.

This work focuses on the problem of autonomous driving in urban environments with unsignalized intersections and in the presence of pedestrians. A deep RL agent is designed for high-level behavioral decision making combined with a classical PID controller for low-level actuation. The agent is trained in a simulator to learn the desired driving policy. This proposed approach is compared with a rule-based one for performance evaluation.

The main contributions of this work are as follows:

- To learn high-level behavioural decision making, a Deep Recurrent Q-Network (DRQN) [4] is developed. The DRQN uses a Long Short-Term Memory (LSTM) layer at the end of a typical Deep Q-Network (DQN) [1] giving memory to deep RL agent.
- A unique multi-channeled image like representation is used to train the deep RL agent. This representation can store a variety of information (geometry as well as semantic) necessary for training in complex urban environments.
- Instead of assuming an ideal low level controller, this work combines a more realistic PID control for longitudinal actuation. For this, appropriate information about the agents state is passed to the network so that the decision maker learns the desired policy taking into account the behavior of the PID controller.

The remainder of this paper is structured as follows: A description of deep reinforcement learning algorithms is mentioned in Section 2 and related work is discussed in Section 3. Section 4 then introduces the complete framework and the proposed deep RL based behavioral decision maker method. In Section 5, details about the training process, the simulation setup and performance evaluation results are mentioned. The final section is assigned to concluding remarks and future work discussion.

2 BACKGROUND

2.1 Reinforcement Learning

In reinforcement learning (RL), at each time step t , the agent receives an observation of the current state s_t of the environment and then it decides to take an action a_t based on a policy π . After the action is performed in the environment, the agent receives a numerical reward r_t and the next state s_{t+1} of the environment. The goal of an RL agent is to learn an optimal policy, which maximizes the sum of future discounted rewards R , defined as :

$$R = \sum_{t=0}^T \gamma \cdot r_t$$

where γ is a discount factor defining the importance of future rewards.

2.2 Q-Learning

In Q-learning, a type of RL algorithm, the agent tries to learn a state-action value function, also called the Q-function. This function for a policy π is defined as the expected return for executing an action a in a state s :

$$Q^\pi(s, a) = E[R_t | s_t = s, a_t = a] \quad (1)$$

The function gives a Q-value for each state-action pair indicating how good taking an action might be at a particular state. The goal of Q-learning algorithm is to find an optimal value function $Q^*(s, a)$ defined as:

$$Q^*(s, a) = \max_{\pi} E[R_t | s_t = s, a_t = a] = \max_{\pi} Q^\pi(s, a) \quad (2)$$

2.3 Deep Q-Networks

In Deep Q-Networks (DQN) [1], a multi-layered neural network, with parameters θ , is used to approximate the value function and is denoted as $Q(s, a; \theta)$. The DQN outputs a vector of action values, given a state s .

Two important ingredients of DQN are [18]: experience replay and a target network. Replay memory improves the performance of learning by eliminating the correlations in the observation sequences. The target network, with parameters θ^- , is same as the main network, however, its parameters are updated from the main network after every t steps, that is, $\theta_t^- = \theta_t$. The goal is to minimize the average of the loss:

$$L(\theta_t) = \left(\underbrace{(r + \gamma \max_a Q(s_{t+1}, a; \theta_t^-) - Q(s, a; \theta_t^-))}_y \right)^2 \quad (3)$$

where t is the time step and y is the update target used by DQN.

2.4 Double Deep Q-Networks

Conventional DQN faces the problem of overestimation. The max operation in DQN uses the same Q-function network for both, action selection and evaluation, causing overestimation of action values. This is solved by using Double Deep Q-Network (DDQN) which uses the main network for evaluating the greedy policy, while the target network is used to estimate the value of this policy. The target for loss function in DDQN, is defined as:

$$y^{DDQN} = r + \gamma Q(s_{t+1}, \underset{a}{\operatorname{argmax}} Q(s_{t+1}, a; \theta_t); \theta_t^-) \quad (4)$$

In the *argmax* term, the main network with weights θ_t , is used for action selection, while the target network with weights θ_t^- is for action evaluation.

2.5 Long Short-Term Memory (LSTM)

A regular recurrent neural network has vanishing gradients problem while dealing with longer sequences. This problem is solved by using an LSTM [19] which stores information in memory cells, carefully regulated by using input, output and forget gates. These gates can learn to decide if a memory cell is to be kept or removed. This helps in storing both recent and older relevant observations to be utilized by the network.

3 RELATED WORK

Given the recent developments in deep learning, interest in using learning based approaches for autonomous driving has increased significantly. There has also been immense interest in using deep RL for autonomous navigation given its ability to self learn from experiences. A detailed survey of these approaches could be found in [5, 6].

One of the important aspect for deep RL is way in which the state is represented and therefore, many different state representations are being used. In [7], [8] raw RGB images are used as a state representation for the RL agent. Another common practice for state representation is the use of a one-dimensional vector storing relative position, speed and lane of the surrounding vehicles. For instance, [9] uses information of 8 surrounding vehicles in the form of a vector for the problem of speed and lane change decision making on highway environment. An image-like grid (matrix) format for state representation has been used in some approaches [10], [11] for dealing with intersection crossing scenarios. In [10] grid cells are used for representing the speed of surrounding vehicles for lane change decision making. The traffic information is extracted from the simulator to create a lane based grid representation. In [11], a multi-layered grid representation is used for urban traffic light control. It consists of two matrices of size 64×64 , one for surrounding vehicle position and another for their speeds.

Raw RGB images contain very high dimensional data such as textures and appearance of roads and objects. Hence for generalization, large data set is required covering each of these dimensions. The vector based representation works well for highway environments where the road geometry and the traffic is structured and the agents behavior is affected by leading and preceding vehicles in the same and opposite lanes. However, urban environments have complex road structures such as intersections and crossings. It consists of different vulnerable road users (VRUs) such as pedestrians and bicyclists with motions that are less constraint in comparison to the vehicles. Hence, in this work an image like multi-layered grid representation state representation is employed due to its capability to represent a variety of features essential for decision making in urban environment context.

One problem when using reinforcement learning for autonomous driving is the assumption that the environment is fully observable, that is, there is complete information about the state of the environment at any step [12]. This problem is solved by some authors, for example, in [13], an approach called Deep Recurrent Q-Network (DRQN) is developed which uses a Long Short-Term memory (LSTM) layer to estimate the Q-function. However, these approaches use feature vector state representation for the DRQN decision maker. Hence, this work takes inspirations from [4] which

uses a similar DRQN approach to train an agent play Atari 2600 games with partial observability. However, in [4] the input to the agent is a multi-channel raw image which passes through several convolutional neural layers and the output of these layers is then passed to an LSTM layer which selecting appropriate action.

4 PROBLEM FORMULATION

Autonomous driving in urban environments is challenging due to the complexity of the road structure and the uncertainty in the behavior of diverse road users. Interpreting the intentions and motions of other road users becomes important, especially for vulnerable road users (VRUs) such as pedestrians as their motions are less constraint in comparison to vehicles and even slightest collision is likely to be fatal. Often, pedestrians perform jaw-walking putting themselves into a vulnerable situation in the context of autonomous driving.

This work focuses on the high-level behavioural decision making aspect of autonomous driving in the presence of pedestrians for a crowded urban environment with unsignalized intersections. A framework is proposed to reduce the complexity of the task by dividing it into sub-problems: high-level decision making and low-level control. For decision making, a Deep Recurrent Q-Network (DRQN) is used and combined with a PID longitudinal controller. The objective is to autonomously drive along an urban road which has unsignalized intersections with crosswalks and control the velocity such that collision with pedestrians could be avoided in a comfortable manner.

4.1 System architecture

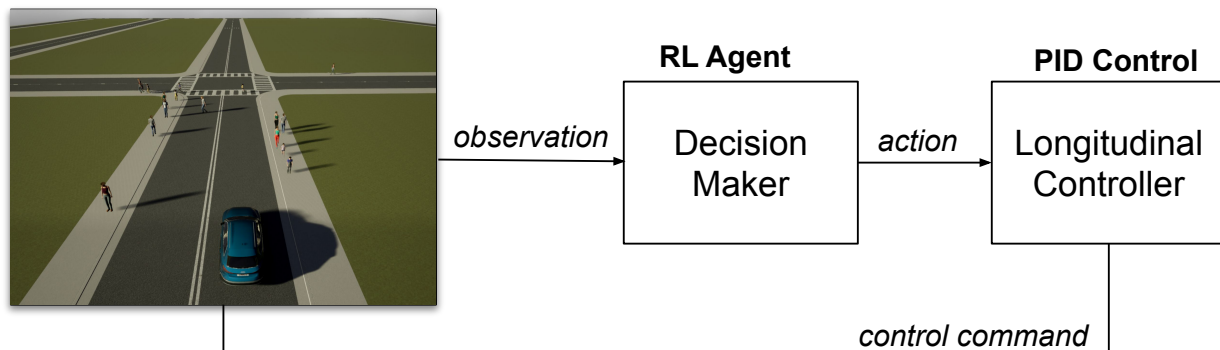


Figure 1: Framework architecture.

Fig. 1 shows the proposed framework architecture. An urban environment is defined as the world around the ego vehicle, including all pedestrians of interest. The environment is developed in a simulator explained in Section 5.4 and sends observations to the agent at each instance. Based on the learned decision making policy, the agent chooses appropriate high level actions (details in Subsec. 4.2.2). These actions are then sent to a longitudinal PID controller which calculates appropriate throttle values given to the ego vehicle.

4.2 POMDP formulation

This work formulates the autonomous driving decision making in urban environment as a Partially Observable Markov Decision Process (POMDP), as the ego-vehicle cannot observe the intentions, that is, the intended destination of the surrounding pedestrians. The input state representation, action space and the reward function for learning the desired driving policy are mentioned below:

4.2.1 State space

The representation for state space needs to have enough and useful information about the environment and the agent itself for urban autonomous driving. Hence, a multi-layered (3-D) grid representation, derived from our previous work [17] is used. It resembles an image-like representation with multiple channels, with each channel corresponding to a particular layer. To generate this grid, a bird's eye view of environment is discretized into a grid in Cartesian coordinates. The resolution of each grid is of size $1m \times 1m$, assuming the footprint of a pedestrian is approximately 1 meter .

The first layer is a position matrix, where a grid value is set to 1 indicating presence of pedestrian and absence of pedestrian is indicated by 0 value. The second layer stores relative heading information about the corresponding pedestrian, and the third layer is use for relative speed. The fourth layer stores semantic information in the form of type of road structure occupied by the pedestrian, namely: walkway, crossing or road.

4.2.2 Action Space

The goal is to learn an optimal decision making policy to select appropriate high-level driving behavior. For this 4 actions are defined, each corresponding to a specific behavior, namely:

- **accelerate**: to speed up by gradually increasing the current speed by $+1 \text{ km/hr}$.
- **slow down**: to slow down by gradually decreasing agents current speed by -1 km/hr .
- **brake**: to apply brakes and stop.
- **steer**: to advance by maintaining the current speed.




Figure 2: Longitudinal PID Controller [20]

The selected action (behavior) is then passed to a PID controller, shown in Fig. 2, which in turn generates appropriate throttle and brake values. This is explained in Algo. 1. Basically, each action updates the desired speed (v_d) input to the controller. Then based on the error between current ego vehicle speed (v) and the updated desired speed (v_d), the controller generates appropriate acceleration (u) using the equation:

$$u = K_p(v_d - v) + \int_0^t (v_d - v)dt + \frac{d(v_d - v)}{dt} \quad (5)$$

The acceleration value further decides the throttle position T_p and braking B_p , using the following conditions:

$$T_p = \begin{cases} u, & \text{if } u \geq 1 \\ 0, & \text{if } u < 0 \end{cases}, \quad B_p = \begin{cases} 0, & \text{if } u \geq 1 \\ -u, & \text{if } u < 0 \end{cases} \quad (6)$$

4.2.3 Reward Function

A reward function helps an agent in learning its desired policy. In the context of autonomous driving in urban environment, two main aspects are considered while developing the reward function: safety and efficiency. With respect to safety, an autonomous vehicle should be able to avoid any collisions while driving. It also must do this without causing discomfort to the pedestrians and its passengers. For efficiency, the autonomous vehicle should drive as fast as possible but without exceeding the speed limits. After evaluating several reward combinations, the following function is used to formulate the reward:

$$r = \begin{cases} r_c & \text{for collision} \\ r_{nc} & \text{for near collision} \\ r_v & \text{for normal driving} \end{cases} \quad (7)$$

where $r_c = -10$, is a large penalty if there is a collision. A small penalty of r_{nc} is also given in case of near collision situations, to discourage the agent from driving too close to the pedestrians and causing them discomfort. To compute r_{nc} , a time-to-collision (TTC) based piece-wise function is used. Pedestrians with $TTC \leq 3$ seconds are considered and the r_{nc} is updated using the following equation:

$$r_{nc} = TTC - 3.0 \quad (8)$$

Algorithm 1: High-level action to throttle, brake conversion

Input: $Actions = \{accelerate, slow_down, brake, steer\}$ **Output:** $Throttle(T_p), Brake(B_p)$ $T_p, B_p \leftarrow 0$ **if acceleration then** $v_d \leftarrow +1$ $T_p, B_p \leftarrow compute_throttle_brake(v_d)$ **end****else if slow_down then** $v_d \leftarrow -1$ $T_p, B_p \leftarrow compute_throttle_brake(v_d)$ **end****else if steer then** $v_d \leftarrow +0$ $T_p, B_p \leftarrow compute_throttle_brake(v_d)$ **end****else if brake then** $T_p \leftarrow 0$ $B_p \leftarrow 1$ **end****return** T_p, B_p

Hence, lesser the time to collision, higher will be the penalty r_{nc} . The term r_v represents a reward for encouraging the agent to maintain a higher speed within the traffic speed limits. The reward for speed is computed using the following function:

$$r_v = \begin{cases} 1 - \lambda \cdot (v_{ref} - v_{ev}), & \text{if } 0 < v_{ev} \leq v_{ref} \\ -1, & \text{if } v_{ev} \leq 0 \\ -0.5, & \text{if } v_{ev} > v_{ref} \end{cases} \quad (9)$$

where $\lambda = 1/v_{ref}$, v_{ev} is ego vehicle speed and v_{ref} is the maximum allowed speed. A penalty of -1 is assigned to the ego vehicle to encourage it to keep moving and a penalty of -0.5 is applied for exceeding the speed limit.

5 EXPERIMENTS AND RESULTS

5.1 Neural Network Model




Figure 3: The network architecture

The network architecture used in this work consists of stacked LSTM structure inspired from [16] and is shown in Fig. 3. The input x_1 to the network is a multi-layered grid of size $45 \times 30 \times 4$. This input is propagated through three convolutional layers with each convolutional layer followed by a ReLU activation function. The third convolutional layer is then flattened and is fed into an LSTM layer. The output of this layer is then input to the following LSTM layer along with input x_2 which consists of ego vehicle's speed and the selected action. Finally the output of the stacked LSTM is fed to the output layer through a fully connected layer. The parameters of the network are mentioned in the Tab. 1 below:

Layer	Activation	Kernel size	Stride	Filters	Units
Conv1	ReLU	8*6	4	32	-
Conv2	ReLU	4*3	3	64	-
Conv3	ReLU	2*2	2	64	-
LSTM1	-	-	-	-	256
LSTM2	-	-	-	-	256
FC	ReLU	-	-	-	256

Table 1: Parameters of the neural network

5.2 Training

During training, the agent interacts with the environment following an ϵ -greedy policy mentioned in section 5.3. The network is trained using Double Deep Q-Network (DDQN) algorithm, described in Sec. 2.4, which consists of one main Q network and a target Q network with the same architecture. The parameters of the main network are updated using an Adam optimizer [14]. The parameters of the target network are updated periodically after every N steps, by cloning the main network parameters. For stability, error clipping is employed by limiting the error term of Eq. (3) to $[-1, 1]$. Since the network uses a recurrent layer, sequential updates are required. The experiences gathered during training are stored as episodes in a replay memory. During training a mini-batch of these episodes is selected uniformly, with each episode consisting of 8 time steps. After performing several tests, a relatively good set of hyper-parameters have been selected and are mentioned in the Tab. 2.

Hyper-parameters	Value	Description
Learning rate, α	0.001	Learning rate used by Adam
Discount factor, γ	0.9	Discounted future reward
Mini-batch size, M_{mini}	32	Number of training episodes
Target update, N	10000	Target network update steps
Replay memory size, M	50	Max. episodes stored in memory
Training episodes, E	200	Total episodes used for training
Episode steps, S	1000	Maximum steps per episode

Table 2: Hyper parameters of the DDQN agent

5.3 Exploration Strategy

It is a common practice to use an ϵ -greedy policy for Q-learning. Generally, epsilon greedy approach randomly selects an action from a uniform distribution. A random action is selected with probability ϵ and current optimal action is selected with probability of $1 - \epsilon$. When $\epsilon = 1$, completely random action is selected, whereas, when $\epsilon = 0$ current optimal action is selected. In this work, the value of ϵ is linearly decreased from $\epsilon_{start} = 1.0$ to $\epsilon_{end} = 0.1$ over the training episodes. However, instead of a uniform distribution, this work employs a modification to improve the learning process. For initial few episodes, the *accelerate* and *steer* actions are selected with slightly higher probability, to encourage the ego vehicle to learn moving at an early stage and subsequently encounter collision situations sooner.

5.4 Simulation setup and scenario

The proposed method is trained and evaluated in CARLA simulator [15], an open source simulator for autonomous driving. An urban environment with unsignalized intersection is developed in the simulator. The simulator provides environment information and receives throttle and brake signals for the ego vehicle. The simulator time step is set at $t = 0.1$ s.

A typical urban environment scenario is considered as shown in Fig. 4. At the beginning of each episode, an ego vehicle is spawned at the start position and follows a list of way-points generated by the simulator to reach its destination. The task for the ego vehicle is to navigate as fast as possible, within the speed limits, through the way-points to reach its destination while comfortably avoiding collisions with the surrounding pedestrians. There are 10 pedestrians initially sampled randomly in the vicinity of the ego vehicle, with random desired speed ranging between $0.5 - 1.5$ km/hr and a randomly assigned goal. Each pedestrian follows one of the three behaviors: legal crossing on crosswalks, illegal jaw walking on roads or normal walking on sidewalks (no crossing) with the probabilities of 0.6, 0.2, 0.2 respectively. As the ego vehicle moves through its route, pedestrians moving farther away are removed and corresponding number of new pedestrians are sampled in the vicinity of the vehicle, making sure there is always quite busy pedestrian traffic.




Figure 4: Simulation environment and training scenario




Figure 5: Training performance

5.5 Results

The proposed DRQN agent is trained for 200 episodes. For each episode the maximum step limit is set to 1000 steps, however, an episode terminates if the goal is reached or if there is a collision. The average reward and the number of steps executed per episode are shown in Fig. 5. As seen in the reward graph, in the beginning, the reward gradually increases as the agent learns to move. However, there is a sudden drop in the reward, because, as the agent learns to move it starts encountering collisions and getting penalties. After some episodes, the agent starts learning the desired behavior and hence the reward starts increasing again. Similarly, initially the number of steps per episode are high as the agent is learning to move. After some episodes, there is a drop in number of steps due to collisions and as training progresses further the number of steps start increasing again.

The performance of the trained agent is evaluated by comparing it with a rule-based method of CARLA simulator. This method always tries to steer the ego vehicle with maximum speed to reach the target destination while avoiding collisions with other vehicles and pedestrians. This is done by applying emergency braking in case of a collision situation. The maximum speed (v_{max}) is set to 15 km/hr , by fixing the desired speed of the PID controller to $v_d = v_{max}$. The rule-based agent follows a conservative policy, wherein, the ego-vehicle starts applying braking for legal and illegal crossing pedestrians which are in the vicinity of upto 7 meters , in front of the vehicle.

Both the methods are tested for 10 episodes each and are evaluated using the metrics of percentage of collision free episodes, average speed (in km/hr) and average distance travelled (in meters). The results are depicted in Tab. 3. In terms of safety, the proposed method performs much better than a standard rule-based one. Also, the average distance travelled without any collisions is much better with the proposed method. In terms of average speed, the rule-based method performs better. This is because, since the DRQN has memory, it anticipates the behavior of nearby pedestrians and slows down in case of a possible collision like situations. The collision and speed performance of the proposed DRQN method could be further enhanced by training the agent on more episodes.

	Collision free episodes	Average speed	Distance travelled
rule-based agent	40%	7.79	82.6
DQN-based agent	70%	6.09	123.1

Table 3: Comparison between rule-based and DQN approaches

6 CONCLUSION

In this paper, a method is proposed to handle the problem of high-level behavioral decision making in autonomous driving amongst pedestrians. For that an urban environment with unsignalized intersections is developed in a real-world-like simulation environment. A Deep Recurrent Q-Network (DRQN) based agent is used for decision making combined with a low level PID controller. Given a multi-layered state representation and a well designed reward function, the trained agent is able to take appropriate actions amongst pedestrians. The trained agent is compared with a rule-based method and the results demonstrate the effectiveness of the proposed method. Although the proposed method is better than rule-based, its performance could further be improved. In the future, the agent will be trained

on a larger dataset using more computational resources to push the limit of the proposed method and also to further improve the learning efficiency. More advanced deep RL algorithms could be explored for behavioral decision making for urban autonomous driving in the future. Another potential future work is to consider more complex urban scenarios by incorporating vehicles as well, along with pedestrians.

7 ACKNOWLEDGMENT

This work was funded under project CAMPUS (Connected Automated Mobility Platform for Urban Sustainability) sponsored by Programme d'Investissements d'Avenir (PIA) of french Agence de l'Environnement et de la Maîtrise de l'Énergie (ADEME).

References

- [1] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. A. Riedmiller, "Playing Atari with Deep Reinforcement Learning", *arXiv preprint arXiv:1509.02971*, pp. 1312-5602, 2013.
- [2] T. Lillicrap, J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning", *arXiv preprint arXiv:1509.02971*, 2015.
- [3] J. Xin, H. Zhao, D. Liu and M. Li, "Application of deep reinforcement learning in mobile robot path planning", In *Chinese Automation Congress (CAC)*, pp. 7112-7116, 2017.
- [4] M. Hausknecht and P. Stone, "Deep Recurrent Q-Learning for Partially Observable MDPs", *arXiv preprint arXiv:1507.06527*, 2015.
- [5] M. Veres and M. Moussa, "Deep Learning for Intelligent Transportation Systems: A Survey of Emerging Trends", In *IEEE Transactions on Intelligent Transportation Systems (TITS)*, pp. 1-17, 2019.
- [6] B. R. Kiran, I. Sobh, V. Talpaert, P. Mannion, A. A. Al Sallab, S. Yogamani and P. Pérez, "Deep Reinforcement Learning for Autonomous Driving: A Survey", *arXiv preprint arXiv:2002.00444*, 2020.
- [7] H. An and J.-i. Jung, "Decision-Making System for Lane Change Using Deep Reinforcement Learning in Connected and Automated Driving", *Electronics*, vol. 8, no. 5, pp. 543, 2019.
- [8] P. Wolf and C. Hubschneider and M. Weber and A. Bauer and J. Härtl and F. Dürr and J. M. Zöllner, "Learning how to drive in a real world simulation with deep Q-Networks", In *IEEE Intelligent Vehicles Symposium (IV)*, pp. 244-250, 2017.
- [9] C. Hoel, K. Wolff and L. Laine, "Automated speed and lane change decision making using deep reinforcement learning", In *IEEE International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2148-2155, 2018.
- [10] J. Wang, Q. Zhang, D. Zhao and Y. Chen, "Lane Change Decision-making through Deep Reinforcement Learning with Rule-based Constraints", In *International Joint Conference on Neural Networks (IJCNN)*, pp. 1-6, 2019.
- [11] J. J. A. Calvo and I. Dusparic, "Heterogeneous Multi-Agent Deep Reinforcement Learning for Traffic Lights Control", In *Irish Conference on Artificial Intelligence and Cognitive Science*, pp. 1-12, 2018.
- [12] S. S. Shwartz, S. Shammah and Amnon Shashua, "Safe, Multi-Agent, Reinforcement Learning for Autonomous Driving", *arXiv preprint arXiv:1610.03295*, 2016.
- [13] P. Wang and C. Y. Chan, "Formulation of deep reinforcement learning architecture toward autonomous driving for on-ramp merge", In *International Conference on Intelligent Transportation Systems (ITSC)*, 2017.
- [14] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization", *arXiv preprint arXiv:1412.6980*, 2014.
- [15] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez and V. Koltun, "CARLA: An Open Urban Driving Simulator", In *Proceedings of the 1st Annual Conference on Robot Learning (CoRL)*, pp. 1-16, 2017.
- [16] M. Liulong, L. Yanjie, C. Jiao, J. Dong, "Learning to Navigate in Indoor Environments: from Memorizing to Reasoning", *arXiv preprint arXiv:1904.06933*, 2019.
- [17] N. Deshpande and A. Spalanzani, "Deep Reinforcement Learning based Vehicle Navigation amongst pedestrians using a Grid-based state representation", In *Intelligent Transportation Systems Conference (ITSC)*, pp. 2081-2086, 2019.
- [18] V. Mnih et al., "Human-level control through deep reinforcement learning", vol. 518, no. 7540, pp.529, 2015.
- [19] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory", In *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997.
- [20] S. Waslander and J. Kelly, "Introduction to Self-Driving Cars", [week 5: Vehicle Longitudinal Control], <https://www.coursera.org/learn/intro-self-driving-cars>.