

Integrating Deep Reinforcement Learning with Model-based Path Planners for Automated Driving

Ekim Yurtsever^{*,†}, Linda Capito^{*,†}, Keith Redmill^{*} and Umit Ozguner^{*}

Abstract—Automated driving in urban settings is challenging. Human participant behavior is difficult to model, and conventional, rule-based Automated Driving Systems (ADSs) tend to fail when they face unmodeled dynamics. On the other hand, the more recent, end-to-end Deep Reinforcement Learning (DRL) based model-free ADSs have shown promising results. However, pure learning-based approaches lack the hard-coded safety measures of model-based controllers. Here we propose a hybrid approach for integrating a path planning pipe into a vision based DRL framework to alleviate the shortcomings of both worlds. In summary, the DRL agent is trained to follow the path planner’s waypoints as close as possible. The agent learns this policy by interacting with the environment. The reward function contains two major terms: the penalty of straying away from the path planner and the penalty of having a collision. The latter has precedence in the form of having a significantly greater numerical value. Experimental results show that the proposed method can plan its path and navigate between randomly chosen origin-destination points in CARLA, a dynamic urban simulation environment. Our code is open-source and available online ¹.

I. INTRODUCTION

Automated Driving Systems (ADSs) promise a decisive answer to the ever-increasing transportation demands. However, widespread deployment is not on the horizon as state-of-the-art is not robust enough for urban driving. The recent Uber accident [1] is an unfortunate precursor: the technology is not ready yet.

There are two common ADS design choices [2]. The first one is the more conventional, model-based, modular pipeline approach [3]–[10]. A typical pipe starts with a perception module. Robustness of perception modules has been increased greatly due to the recent advent of deep Convolutional Neural Networks (CNN) [11]. The pipe usually continues with scene understanding [12], assessment [13], planning [14] and finally ends with motor control. The major shortcomings of modular model-based planners can be summarized as complexity, error propagation, and lack of generalization outside pre-postulated model dynamics.

The alternative end-to-end approaches [15]–[24] eliminated the complexity of conventional modular systems. With the recent developments in the machine learning field, sensory inputs now can directly be mapped to an action space. Deep Reinforcement Learning (DRL) based frameworks can learn to drive from front-facing monocular camera images




Fig. 1. An overview of our framework. FC stands for Fully Connected layers. The proposed system is a hybrid of a model-based planner and a model-free DRL agent. *Other sensor inputs can be anything the conventional pipe needs. ** We integrate planning into the DRL agent by adding ‘distance to the closest waypoint’ into our state-space, where the path planner gives the closest waypoint. Any kind of path planner can be integrated into the DRL agent with the proposed method.

directly [21]. However, the lack of hard-coded safety measures, interpretability, and direct control over path constraints limit the usefulness of these methods.

We propose a hybrid methodology to mitigate the drawbacks of both approaches. In summary, the proposed method integrates a short pipeline of localization and path planning modules into a DRL driving agent. The training goal is to teach the DRL agent to oversee the planner and follow it if it is safe to follow. The proposed method was implemented with a Deep Q Network (DQN) [25] based RL agent and the A* [26] path planner. First, the localization module outputs the ego-vehicle position. With a given destination point, the path planner uses the A* algorithm [26] to generate a set of waypoints. The distance to the closest waypoint, along with monocular camera images and ego-vehicle dynamics, are then fed into the DQN based RL agent to select discretized steering and acceleration actions. During training, the driving agent is penalized for making collisions and being far from the closest waypoint asymmetrically, with the former term having precedence. We believe this can make the agent prone to follow waypoints during free driving but have enough flexibility to stray from the path for collision avoidance using visual cues. An overview of the proposed approach is shown in Figure 1.

The major contributions of this work can be summarized as follows:

- A general framework for integrating path planners into model-free DRL based driving agents
- Implementation of the proposed method with an A* planner and a DQN RL agent. Our code is open-source and available online¹.

The remainder of the paper is organized in five sections.

^{*}E. Yurtsever, L. Capito, K. Redmill and U. Ozguner are with The Ohio State University, Ohio, US.

[†] These authors contributed equally to this work

Corresponding author: Ekim Yurtsever, ekimyurtsever@gmail.com

¹<https://github.com/Ekim-Yurtsever/Hybrid-DeepRL-Automated-Driving>

A brief literature survey is given in Section II. Section III explains the proposed methodology and is followed by experimental details in Section IV. Results are discussed in Section V and a short conclusion is given in Section VI.

II. RELATED WORKS

End-to-end driving systems use a single algorithm/module to map sensory inputs to an action space. ALVINN [16] was the first end-to-end driving system and utilized a shallow, fully connected neural network to map image and laser range inputs to a discretized direction space. The network was trained in a supervised fashion with labeled simulation data. More recent studies employed real-world driving data and used convolutional layers to increase performance [18]. However, real-world urban driving has not been realized with an end-to-end system yet.

A CNN based partial end-to-end system was introduced to map the image space to a finite set of intermediary “affordance indicators” [15]. A simple controller logic was then used to generate driving actions from these affordance indicators. Chauffer Net [27] is another example of a mid-to-mid system. These systems benefit from robust perception modules on the one end, and rule-based controllers with hard-coded safety measures on the other end.

All the methods mentioned above suffer from shortcomings of supervised learning—namely, a significant dependency on labeled data, overfitting, and lack of interpretability. Deep Reinforcement Learning (DRL) based automated driving agents [20], [21] replaced the need for huge amounts of labeled data with online interaction. DRL agents try to learn the optimum way of driving instead of imitating a target human driver. However, the need for interaction raises a significant issue. Since failures cannot be tolerated for safety-critical applications, in almost all cases, the agent must be trained in a virtual environment. This adds the additional virtual-to-real transfer learning problem to the task. In addition, DRL still suffers from a lack of interpretability and hard-coded safety measures.

A very recent study [28] focused on general tactical decision making for automated driving using the AlphaGo Zero algorithm [29]. AlphaGo Zero combines tree-search with neural networks in a reinforcement learning framework, and its implementation to the automated driving domain is promising. However, this study [28] was limited to only high-level tactical driving actions such as staying on a lane or making a lane change.

Against this backdrop, here we propose a hybrid DRL-based driving automation framework. The primary motivation is to integrate path-planning into DRL frameworks for achieving a more robust driving experience and a faster learning process.

III. PROPOSED METHOD

A. Problem formulation

In this study, automated driving is defined as a Markov Decision Process (MDP) with the tuple of (S, A, P, r) . We




Fig. 2. Illustration of state $s_t \simeq (z_t, e_t, d_t)$ and distance to the final destination l_t at time t . Waypoints $w \in W$ are to be obtained from the path planner.

integrate path-planning into the MDP by adding d , distance to the closest waypoint, to the state-space.

- S : A set of states. We associate observations made at time t with the state s_t as $s_t \simeq (z_t, e_t, d_t)$ where; 1) $z_t = f_{cnn}(I_t)$ is a visual feature vector which is extracted using a deep CNN from a single image I_t captured by a front-facing monocular camera. 2) e_t is a vector of ego-vehicle states including speed and location 3) d_t is the distance to the closest waypoint obtained from the model-based path planner. d_t is the key observation which links model-based path planners to the MDP.
- A : A set of discrete driving actions illustrated in Figure 3. Actions consist of discretized steering angle and acceleration values. The agent executes actions to change states.
- P : The transition probability $P_t = \Pr(s_{t+1}|s_t, a_t)$. Which is the probability of reaching state s_{t+1} after executing action a_t in state s_t .
- r : A reward function $r(s_{t+1}, s_t, a_t)$. Which gives the instant reward of going from state s_t to s_{t+1} with a_t .

The goal is to find a policy function $\pi(s_t) = a_t$ that will select an action given a state such that it will maximize the following expectation of cumulative future rewards where s_{t+1} is taken from P_t .

$$\mathbb{E} \left(\sum_{t=0}^{\infty} \gamma^t r(s_t, s_{t+1}, a_t) \right) \quad (1)$$

Where γ is the discount factor, which is a scalar between $0 \leq \gamma \leq 1$ that determines the relative importance of later rewards with respect to previous rewards. We fix the horizon for this expectation with a finite value in practice.

Our problem formulation is similar to a previous study [21], the critical difference being the addition of d_t to the




Fig. 3. The DQN based DRL agent. FC stands for fully connected. After training, the agent selects the best action by taking the argmax of predicted Q values.

state space and the reward function. An illustration of our formulation is shown in Figure 2.

B. Reinforcement Learning

Reinforcement learning is an umbrella term for a large number of algorithms derived for solving the Markov Decision Problems (MDP) [21].

In our framework, the objective of reinforcement learning is to train a driving agent who can execute ‘good’ actions so that the new state *and* possible state transitions until a finite expectation horizon will yield a high cumulative reward. The overall goal is quite straightforward for driving: not making collisions and reaching the destination should yield a good reward and vice versa. It must be noted that RL frameworks are not greedy unless $\gamma = 0$. In other words, when an action is chosen, not only the immediate reward but the cumulative rewards of all the expected future state transitions are considered.

Here we employ DQN [25] to solve the MDP problem described above. The main idea of DQN is to use neural networks to approximate the optimal action-value function $Q(s, a)$. This Q function maps the state-action space to \mathbb{R} . $Q: S \times A \rightarrow \mathbb{R}$ while maximizing equation 1. The problem comes down to approximate or to learn this Q function. The following loss function is used for Q-learning at iteration i .

$$L_i(\theta) = \mathbb{E}_{(s, a, r)} \left[\left(r + \gamma \max_{a_{t+1}} Q^{\theta_i^-}(s_{t+1}, a_{t+1}) - Q^{\theta_i}(s_t, a_t) \right)^2 \right] \quad (2)$$

Where Q-Learning updates are applied on samples $(s, a, r) \sim U(D)$. $U(D)$ draws random samples from the data batch D . θ_i is the Q-network parameters and θ_i^- is the target network parameters at iteration i . Details of DQN can be found in [25].

C. Integrating path planning into model-free DRL frameworks

The main contribution of this work is the integration of path planning into DRL frameworks. We achieve this by modifying the state-space with the addition of d . Also, the reward function is changed to include a new reward term r_w , which rewards being close to the nearest waypoint obtained from the model-based path planner, i.e. a small d . Utilizing waypoints to evaluate a DRL framework were suggested in a very recent work [30], but their approach does not consider integrating the waypoint generator into the model.

The proposed reward function is as follows.

$$r = \beta_c r_c + \beta_v r_v + \beta_l r_l + \beta_w r_w \quad (3)$$

Where r_c is the no-collision reward, r_v is the not driving very slow reward, r_l is being-close to the destination reward, and r_w is the proposed being-close to the nearest waypoint reward. The distance to the nearest waypoint d is shown in Figure 2. The weights of these rewards, $\beta_c, \beta_v, \beta_l, \beta_w$, are parameters defining the relative importance of rewards. These parameters are determined heuristically. In the special case of $\beta_c = \beta_v = \beta_l = 0$, the integrated model should mimic the model-based planner.

Please note that any planner, from the naive A* to more complicated algorithms with complete obstacle avoidance capabilities, can be integrated into this framework as long as they provide a waypoint.

IV. EXPERIMENTS

As in all RL frameworks, the agent needs to interact with the environment and fail a lot to learn the desired policies. This makes training RL driving agents in real-world extremely challenging as failed attempts cannot be tolerated. As such, we focused only on simulations in this study. Real-world adaptation is outside of the scope of this work.

The proposed method was implemented in Python based on an open-source RL framework [31] and CARLA [32] was used as the simulation environment. The commonly used

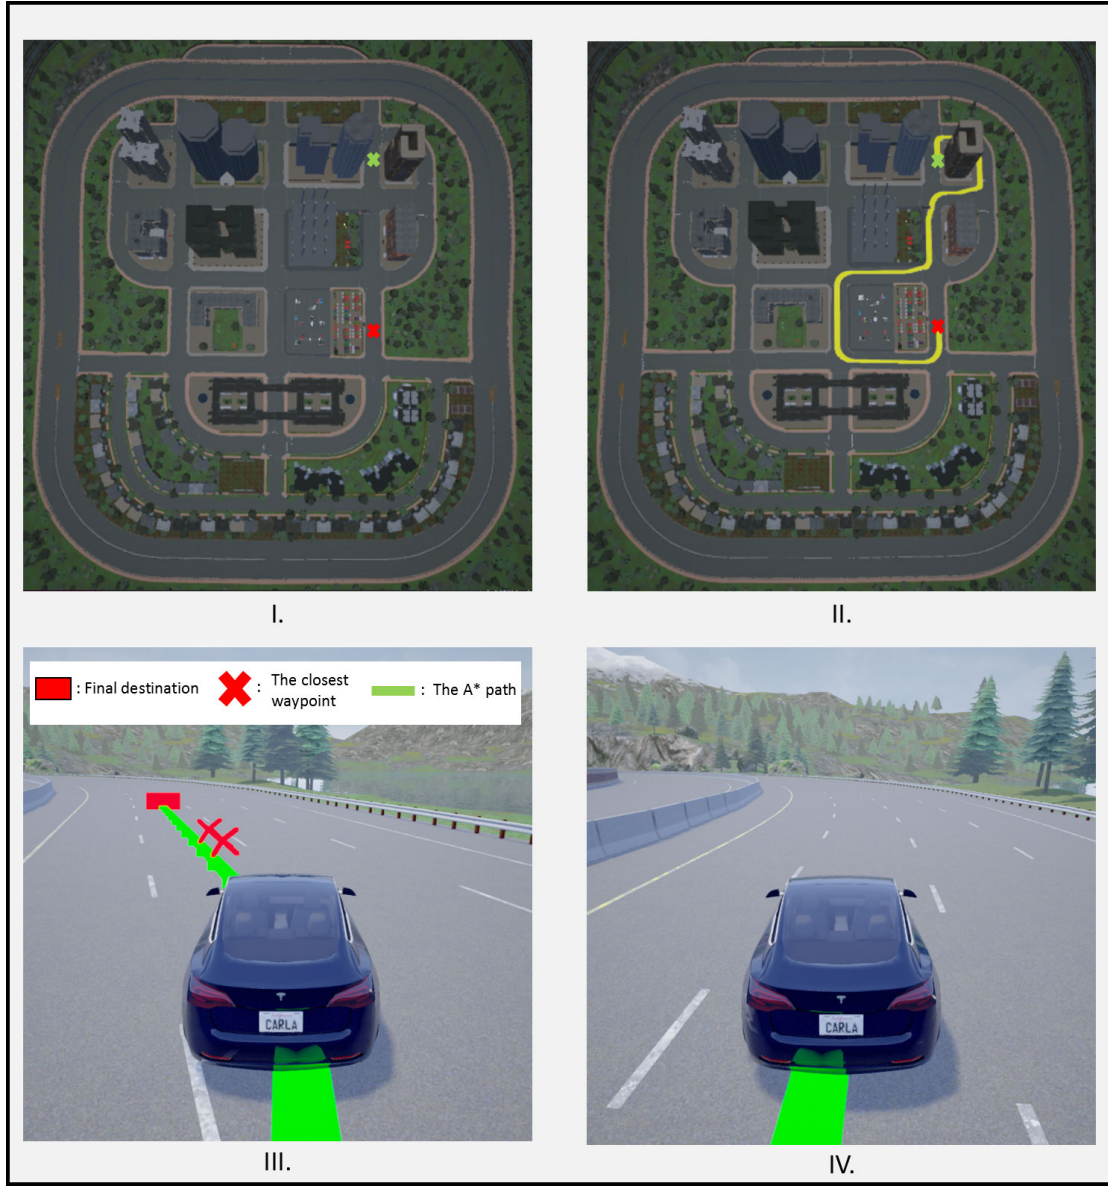


Fig. 4. The experimental process: I. A random origin-destination pair was selected. II. The A* algorithm was used to generate a path. III. The hybrid DRL agent starts to take action with the incoming state stream. IV. The end of the episode.

A* algorithm [26] was employed as the model-based path planner, and the recently proposed DQN [25] was chosen as the model-free DRL.

A. Details of the reward function

The general form of r was given in the previous Section in equation 3. Here, the special case and numerical values used throughout the experiments are explained.

$$r = \begin{cases} r_v + r_l + r_w & , r_c = 0 \text{ \& } l \geq \epsilon \\ 100 & , r_c = 0 \text{ \& } l < \epsilon \\ r_c & , r_c \neq 0 \end{cases} \quad (4)$$

$$r_c = \begin{cases} 0 & , \text{no collision} \\ -1 & , \text{there is a collision} \end{cases} \quad (5)$$

$$r_v = \frac{1}{v_0}v - 1 \quad (6)$$

$$r_l = 1 - \frac{l}{l_{\text{previous}}} \quad (7)$$

$$r_w = 1 - \frac{d}{d_0} \quad (8)$$

Where $\epsilon = 5m$, the desired speed $v_0 = 50km/h$, and $d_0 = 8m$. In summary, r_w rewards keeping a distance less than d_0 to the closest waypoint at every time step, and r_l rewards decreasing l over l_{previous} , distance to the destination in the the previous time step. The last term of r_l allows to continuously penalize/reward the agent for getting further/closer to the final destination.

If there is a collision, the episode is over and the reward gets a penalty equal to -1 . If the vehicle reaches its destination $\exists \epsilon > 0 : l < \epsilon$, a reward of 100 is sent back. Otherwise, the reward consists of the sum of the other terms. d_0 was selected as $8m$ because the average distance between waypoints of the A* equals to this value.

B. DQN architecture and hyperparameters

The deep neural network architecture employed in the DQN is shown in Figure 3. The CNN consisted of three identical convolutional layers with 64 filters and a 3×3 window. Each convolutional layer was followed by average pooling. After flattening, the output of the final convolutional layer, ego-vehicle speed and distance to the closest waypoint were concatenated and fed into a stack of two fully connected layers with 256 hidden units. All but the last layer had rectifier activation functions. The final layer had a linear activation function and outputted the predicted Q values, which were used to choose the optimum action by taking argmax .

Q

C. Experimental process & training

The experimental process is shown in Figure 4. The following steps were carried repeatedly until the agent learned to drive.

- 1) Select two random points on the map as an origin-destination pair for each episode
- 2) Use A* path planner to generate a path between origin-destination using the road topology graph of CARLA.
- 3) Start feeding the stream of states, including distance to the closest waypoint, into the DRL agent. DRL agent starts to take actions at this point. If this is the first episode, initialize the DQN with random weights.
- 4) End the episode if a collision is detected, or the goal is reached.
- 5) Update the weights of the DQN after each episode with the loss function given in equation 2.
- 6) Repeat the above steps sixty thousand times

D. Comparison and evaluation

The proposed hybrid approach was compared against a complete end-to-end DQN agent. The complete end-to-end agent took only monocular camera images and ego-vehicle speed as input. The same network architecture was employed for both methods.

A human driving experiment was also conducted to serve as a baseline. The same reward function that was used to train the DRL agent was used as the evaluation metric. Four adults aging between 25 to 30 years old participated in the experiments. The participants drove a virtual car in CARLA using a keyboard and were told to follow the on-screen path (marked by a green line). The participants did not see their scores. Every participant drove each of the seven predefined routes five times. The average cumulative reward of each route was accepted as the “average human score.”




Fig. 5. Normalized reward versus episode number. The proposed hybrid approach learned to drive faster than its complete end-to-end counterpart.

TABLE I

AVERAGE REWARD SCORES FOR FIVE RUNS IN EACH ROUTE TYPE.

| Route type | Hybrid-DQN | Human average |
|----------------------------|------------|---------------|
| Straight (highway) | 21.1 | 43.4 |
| Straight (urban) | 27.6 | 38.1 |
| Straight (under bridge) | 31.6 | 45.2 |
| Slight curve | 30.4 | 49.5 |
| Sharp curve | -74.4 | -8.9 |
| Right turn in intersection | -136.9 | -12.1 |
| Left turn in intersection | -385.9 | -25.5 |

V. RESULTS

Figure 5 illustrates the training process. The result is clear and evident: The proposed hybrid approach learned to drive much faster than its complete end-to-end counterpart. It should be noted that the proposed approach made a quick jump at the beginning of the training. We believe the waypoints acted as a ‘guide’ and made the algorithm learn faster that way. Our method can be used for spooling up the training process of a complete end-to-end variant with transfer learning. Qualitative analysis of the driving performance can be done by watching the simulation videos on our repository¹.

The proposed method outperformed the end-to-end DQN, however, it is still not good as the average human driver as can be seen in Table I.

Even though promising results were obtained, the experiments at this stage can only be considered as proof of concepts, rather than an exhaustive evaluation. The proposed method needs to consider other integration options, be compared against other state-of-the-art agents, and eventually should be deployed to the real-world and tested there.

The model-based path planner tested here is also very naive. In addition, the obstacle avoidance capabilities of the proposed method was not evaluated. Future experiments should focus on this aspect. The integration of more complete path planners with full obstacle avoidance capabilities can yield better results.

VI. CONCLUSIONS

In this study, a novel hybrid approach for integrating path planning into model-free DRL frameworks was proposed. A proof-of-concept implementation and experiments in a virtual environment showed that the proposed method is capable of learning to drive.

The proposed integration strategy is not limited to path planning. Potentially, the same state-space modification and reward strategy can be applied for integrating vehicle control and trajectory planning modules into model-free DRL agents.

Finally, the current implementation was limited to output only discretized actions. Future work will focus on enabling continuous control and real-world testing.

ACKNOWLEDGMENT

This work was funded by the United States Department of Transportation under award number 69A3551747111 for Mobility21: the National University Transportation Center for Improving Mobility.

Any findings, conclusions, or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the United States Department of Transportation, Carnegie Mellon University, or The Ohio State University.

REFERENCES

- [1] P. Kohli and A. Chadha, "Enabling pedestrian safety using computer vision techniques: A case study of the 2018 uber inc. self-driving car crash," in *Future of Information and Communication Conference*. Springer, 2019, pp. 261–279.
- [2] E. Yurtsever, J. Lambert, A. Carballo, and K. Takeda, "A survey of autonomous driving: Common practices and emerging technologies," *IEEE Access*, vol. 8, pp. 58 443–58 469, 2020.
- [3] C. Urmson, J. Anhalt, D. Bagnell, C. Baker, R. Bittner, M. Clark, J. Dolan, D. Duggins, T. Galatali, C. Geyer, et al., "Autonomous driving in urban environments: Boss and the urban challenge," *Journal of Field Robotics*, vol. 25, no. 8, pp. 425–466, 2008.
- [4] J. Levinson, J. Askeland, J. Becker, J. Dolson, D. Held, S. Kammel, J. Z. Kolter, D. Langer, O. Pink, V. Pratt, et al., "Towards fully autonomous driving: Systems and algorithms," in *Intelligent Vehicles Symposium (IV)*, 2011 IEEE. IEEE, 2011, pp. 163–168.
- [5] J. Wei, J. M. Snider, J. Kim, J. M. Dolan, R. Rajkumar, and B. Litkouhi, "Towards a viable autonomous driving research platform," in *Intelligent Vehicles Symposium (IV)*, 2013 IEEE. IEEE, 2013, pp. 763–770.
- [6] A. Broggi, M. Buzzoni, S. Debattisti, P. Grisleri, M. C. Laghi, P. Medici, and P. Versari, "Extensive tests of autonomous driving technologies," *IEEE Transactions on Intelligent Transportation Systems*, vol. 14, no. 3, pp. 1403–1415, 2013.
- [7] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The oxford robotcar dataset," *The International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.
- [8] N. Akai, L. Y. Morales, T. Yamaguchi, E. Takeuchi, Y. Yoshihara, H. Okuda, T. Suzuki, and Y. Ninomiya, "Autonomous driving based on accurate localization using multilayer lidar and dead reckoning," in *IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017, pp. 1–6.
- [9] E. Guizzo, "How google's self-driving car works," *IEEE Spectrum Online*, vol. 18, no. 7, pp. 1132–1141, 2011.
- [10] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, C. G. Keller, et al., "Making Bertha drive – an autonomous journey on a historic route," *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 2, pp. 8–20, 2014.
- [11] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [12] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [13] E. Yurtsever, Y. Liu, J. Lambert, C. Miyajima, E. Takeuchi, K. Takeda, and J. H. Hansen, "Risky action recognition in lane change video clips using deep spatiotemporal networks with segmentation mask transfer," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*. IEEE, 2019, pp. 3100–3107.
- [14] M. McNaughton, C. Urmson, J. M. Dolan, and J.-W. Lee, "Motion planning for autonomous driving with a conformal spatiotemporal lattice," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 4889–4895.
- [15] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, "Deepdriving: Learning affordance for direct perception in autonomous driving," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2722–2730.
- [16] D. A. Pomerleau, "Alvin: An autonomous land vehicle in a neural network," in *Advances in neural information processing systems*, 1989, pp. 305–313.
- [17] U. Muller, J. Ben, E. Cosatto, B. Flepp, and Y. L. Cun, "Off-road obstacle avoidance through end-to-end learning," in *Advances in neural information processing systems*, 2006, pp. 739–746.
- [18] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, et al., "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [19] H. Xu, Y. Gao, F. Yu, and T. Darrell, "End-to-end learning of driving models from large-scale video datasets," *arXiv preprint*, 2017.
- [20] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep reinforcement learning framework for autonomous driving," *Electronic Imaging*, vol. 2017, no. 19, pp. 70–76, 2017.
- [21] A. Kendall, J. Hawke, D. Janz, P. Mazur, D. Reda, J.-M. Allen, V.-D. Lam, A. Bewley, and A. Shah, "Learning to drive in a day," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 8248–8254.
- [22] S. Baluja, "Evolution of an artificial neural network based autonomous land vehicle controller," *IEEE Transactions on Systems, Man, and Cybernetics-Part B: Cybernetics*, vol. 26, no. 3, pp. 450–463, 1996.
- [23] J. Koutnfk, G. Cuccu, J. Schmidhuber, and F. Gomez, "Evolving large-scale neural networks for vision-based reinforcement learning," in *Proceedings of the 15th annual conference on Genetic and evolutionary computation*. ACM, 2013, pp. 1061–1068.
- [24] K. Makantasis, M. Kontorinaki, and I. Nikolas, "A deep reinforcement learning driving policy for autonomous road vehicles," *arXiv preprint arXiv:1905.09046*, 2019.
- [25] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [26] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [27] M. Bansal, A. Krizhevsky, and A. Ogale, "Chauffeurnet: Learning to drive by imitating the best and synthesizing the worst," *arXiv preprint arXiv:1812.03079*, 2018.
- [28] C.-J. Hoel, K. Driggs-Campbell, K. Wolff, L. Laine, and M. Kochenderfer, "Combining planning and deep reinforcement learning in tactical decision making for autonomous driving," *IEEE Transactions on Intelligent Vehicles*, 2019.
- [29] D. Silver, J. Schrittwieser, K. Simonyan, I. Antonoglou, A. Huang, A. Guez, T. Hubert, L. Baker, M. Lai, A. Bolton, et al., "Mastering the game of go without human knowledge," *Nature*, vol. 550, no. 7676, pp. 354–359, 2017.
- [30] B. Osinski, A. Jakubowski, P. Milos, P. Ziecina, C. Galias, and H. Michalewski, "Simulation-based reinforcement learning for real-world autonomous driving," *arXiv preprint arXiv:1911.12905*, 2019.
- [31] Sentdex, "Carla-rl," <https://github.com/Sentdex/Carla-RL>, 2020.
- [32] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "Carla: An open urban driving simulator," *arXiv preprint arXiv:1711.03938*, 2017.