IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⋇ Impact Factor 7.39 ⋇ Vol. 11, Issue 7, July 2022

DOI: 10.17148/IJARCCE.2022.11789

End-to-End Learning in Autonomous Driving

S Advaith[1], Rachita Agarwal[2], Merin Meleet[3]

Student, Information Science and Engineering, RV College of Engineering, Bangalore, India[1]

Student, Information Science and Engineering, RV College of Engineering, Bangalore, India[2]

Assistant Professor, Information Science and Engineering, RV College of Engineering, Bangalore, India[3]

Abstract: Research has shown that erratic human behaviour such as impaired driving, drugged driving, unbelted vehicle occupants, speeding and distraction are factors in as much as 94% of the crashes in roads today. Automation in this field has the potential to enormously reduce the incidence of such crashes. Higher levels of autonomy mean that the problem of driving is no longer one that human drivers have to solve. This is an area that has gained considerable traction in recent years. The modus operandi in the research community is to understand human driving behaviour and build autonomous units that can imitate this behaviour. This complicated issue domain calls for elaborate solutions that frequently involve numerous modules operating in unison. Each of these modules deals with a particular issue and transmits its solutions to the succeeding modules for processing. The vehicle's controller component, which carries out the predetermined behaviour, receives the ultimate outcome. Additionally, since everything must occur in real-time, prediction speed is just as crucial as underlying accuracy.

This sophisticated modular design has been discovered to be inefficient, and deep learning has been found to be a good replacement. Deep learning involves automatically learning complex mathematical functions that characterise a specific domain. Understanding human drivers is a complex task. It involves emotions rather than logic, and these are all fuelled with reactions. It becomes very uncertain what the next action will be of the drivers or pedestrians nearby, so a system that can predict the actions of other road users can be very important for road safety. The car can observe, gather all the information it requires, and interpret it thanks to a 360-degree vision of its surroundings. Once the data is loaded into the learning system, it can think of every possible move those other drivers could make. It resembles a game in which the player must choose the best move from a limited number of options in order to beat the opposition. In this issue area, an autonomous unit's functions include localising the vehicle in its environment, enhancing perception, and actuating kinematic motions in self-driving cars. This guarantees both easy commuting and road safety.

I. INTRODUCTION

Self-driving technology is arguably the most contentious subject in the automotive sector right now, and several businesses are working to create features for production automobiles that would allow for autonomous driving. The goal of this technology competition is to increase vehicle efficiency and safety. More than 1 million people die on the roads each year as a result of automotive accidents, according to the World Health Organization, and the C2ES (Centre for Climate and Energy Solutions) estimates that automobiles account for nearly 60% of all energy used in transportation. These figures demonstrate that automobiles are a significant source of greenhouse gas emissions and serious casualties. Although attempts have been made to address these issues, none of the remedies have proven to be especially successful. Self-driving technology, however, has the power to solve all of these issues. Self-driving cars could save precious lives and protect the environment if they are operated more safely and effectively. The ability of autonomous vehicles (AVs) to provide driverless mobility to a wide variety of individuals from door to door has the potential to make a substantial impact on society. This new technology will have an impact on both those who can currently operate a vehicle and those who can't because of medical, psychological, or other issues. As experts predict that AVs can significantly cut travel time due to less traffic congestion, this new technology has the potential to help society in a variety of ways. Additionally, anyone travelling in the AV who doesn't drive can be productive while doing so because it is a much safer vehicle. The likelihood of an accident occurring will be lower in AVs than in vehicles driven by humans. Public transportation may become obsolete because AVs will be more affordable, safer due to a reduction in accidents, and more comfortable in terms of privacy and hygiene. In fact, autonomous vehicles might even be a way to maintain economic activity in the event of a global pandemic where public transportation is scarcely used any more while lowering the risk of viral transmission.

Artificial neural networks, a class of algorithms inspired by the structure and operation of the brain, are the focus of the machine learning discipline known as deep learning. The machine equivalent of experience, it enables software to learn from data, but in order to be useful, it needs a tremendous amount of computing power. However, recent advancements in the field of GPU-based processing have made the use of Deep Neural Networks possible in a variety of fields. With

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified ⁙ Impact Factor 7.39 ⁙ Vol. 11, Issue 7, July 2022
DOI: 10.17148/IJARCCE.2022.11789

approximately 89% of accidents being due to human errors, autonomous vehicles will play a vital role to significantly reduce this number and ultimately to save human lives. Artificial Intelligence (AI) will play a crucial part in achieving this objective by enabling the task of explicitly establishing rules to create a system that is able to learn those rules. Older and disabled persons will have more mobility thanks to autonomous vehicles, which may also use up to 90% less energy overall. Additionally, decreased traffic congestion and accompanying air pollution will result. Thus, it is crucial to comprehend and replicate the capability of steering a vehicle using visual cues within an AI system. To the authors' knowledge, Pomerleau [1] completed the first work in this area in 1989 and used a multilayer perceptron (MLP). Since then, considerable parallelization in contemporary GPUs has allowed for a significant increase in computational capacity. combined with contemporary network topologies and recently developed ideas like Convolutional Neural Networks (CNN). These developments made it possible to develop systems that use CNN for frame-by-frame classification of driving scenarios to learn how to steer autonomous vehicles.

II. LITERATURE SURVEY

Autonomous cars have been a topic of increasing interest in recent years. More and more research is being done with a focus on developing hardware and software technologies for fully capable autonomous driving units requiring no human intervention. The following section briefly summarizes the salient points of some of the most prominent work in this domain. In [1], the authors proposed a complete software stack for autonomous driving based on a neural network system that could automatically determine which elements in the road images most influence steering decisions. Road tests demonstrated that the model proposed can successfully perform lane keeping in a wide variety of driving conditions, regardless of whether lane markings are present or not. The model learns all the obvious features and the more subtle ones that would be hard to anticipate by human engineers. The exact applicability of deep learning in this domain has been studied by the authors in [2] as they highlight the power of deep learning architectures in terms of reliability and efficient real-time performance. They also overviewed state-of-the-art strategies for safe autonomous driving, with their major achievements and limitations. Furthermore, this survey includes measurement, analysis, and execution, with a focus on road, lane, vehicle, pedestrian & drowsiness detection, collision avoidance, and traffic sign detection through sensing and vision-based deep learning methods. For the purposes of testing deep neural networks for autonomous driving in a safe emulator environment, the authors in [3] collected artificial data and simulated the process of moving a car using the emulator built by the Udacity team. Furthermore, it aimed to check if this approach is effective for rapid data collection, modelling, training and testing. In [4], the authors show that a convolutional network can be designed which uses the real driving data obtained through the vehicle’s camera and other sensors. The response of the driver during driving is recorded in different situations, by converting the real driver’s driving video to images and obstacle detection is carried out with the best accuracy and speed using the ‘Yolo’ algorithm. This way, the network learns the response of the driver to obstacles in different locations. As is the case in any machine learning model, some more so than others, Deep Learning models are susceptible to adversarial attacks during training. In [5], the authors conducted research into the safety implications of these attacks in the self-driving domain that have been shown to be of vital interest. The proposed system in [6] is attempting to find a function that directly maps raw input images to a predicted steering angle as output using a deep neural network. The CNN model parameters were trained by using data collected from a vehicle platform built with a 1/10 scale RC car, Raspberry Pi 3 Model B computer and front-facing camera. The training data were road images paired with the time-synchronized steering angle generated by manually driving. The experimental results demonstrate the effectiveness and robustness of the autopilot model in lane keeping tasks. In [7], it was shown that deep neural networks have the capability to learn highly nonlinear functions. Furthermore, they don’t require explicit hand-generated algorithms for object detection, lane detection and path planning. Testing of networks is a crucial aspect of the model building process as it ensures reliability and robustness. Offline and online forms of testing in the context of autonomous driving systems were studied in [8]. The authors concluded that while online testing is a highly optimal manner of testing, constructing vivid and accurate simulations of the environment is a difficult task.

III. DATA PREPARATION

For the purposes of testing, the CarND Udacity Simulator, built in Unity, was utilized. With this simulator, driving an automobile may be practiced in both manual and autonomous modes. The software is quite adaptable and enables card customization. In manual mode, a few distinct driving styles were simulated over the course of an hour.

Although much effort was made to avoid bumping into things and deviating off the track, it is still possible to model styles that in the actual world would only have limited applicability. After the ride was recorded and processed, a log was created with several modelling factors, the most important of which are simulated views from the left, right, and centre cameras as well as the steering angle.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁘ Impact Factor 7.39 ⁘ Vol. 11, Issue 7, July 2022

DOI: 10.17148/IJARCCE.2022.11789

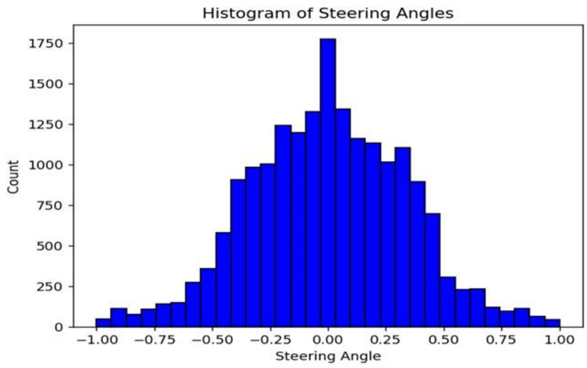

Fig. 1. Histogram depicting distribution of steering angles in Track-2

Roughly 30000 training examples were collected consisting of optimal driving behaviour in the terrain offered by the simulator (track-1). A further 15000 were collected consisting of recovery behaviour in the more complex track-2. This was done for the 'Transfer Learning' process. All data collected is of supervised type with the state of the vehicle (as noted by the corresponding images at that time-step) and the steering angle recorded at that time-step. This makes it possible to perform the recommended 70:30 training-validation split. The use of recovery data permitted the development of significantly better features with integrated safety procedures. The model gains knowledge of the problems of its previous iteration and becomes more adept at navigating sweeping turns.

Fig. 2. Tracks 1,2 & 3

IV. PIPELINE

The state of the vehicle is monitored by way of its image vectors and its speed. Based on the current state of the vehicle and the learned policy, the appropriate control is passed to the controller to execute on behalf of the planning mechanism. In the simulator, this translates to deploying the model and passing on the appropriate actuator control values through an opened socket connection so that the car is driven in that environment. The figure below summarizes the salient features of this approach. More precisely, the state of the car is recorded as the combination of its current speed and 3 images captured at every time-step: front-facing, left-facing and right-facing.

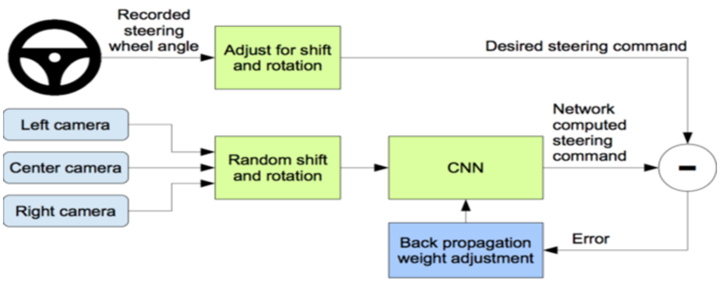

Fig. 3. Internal Design of the system

Training the network as illustrated in the architecture section is done by posing the problem as one of supervised learning. The model is made to learn the optimal driving policy (steering angle) by continuously improving on its errors. This process is known as backpropagation. For the purpose of training, random shifts and transformations are applied on the

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified ⋇ Impact Factor 7.39 ⋇ Vol. 11, Issue 7, July 2022
DOI: 10.17148/IJARCCE.2022.11789

images obtained in the dataset generation process. This process is known as image augmentation. Optimal predictions are then made on the new terrain and passed on to the vehicle on the basis of the learned policy. In the case of the simulator, the controls are passed via socket to execute on the car.

A. Model Architecture

The suggested method makes use of the Nvidia-popularized "Pilot-net" architecture. This is an example of a "Convolutional Neural Network" (or CNN), where various characteristics of an input image (describing the condition of the autonomous vehicle) are given "importance," as seen in the illustration below (represented by learnable weights and biases).

The Visual Cortex served as inspiration for CNN's architecture, which is similar to the connectivity network of neurons in the human brain. Only in this constrained area of the visual field, known as the Receptive Field, do individual neurons react to stimuli. These fields are grouped together and cover the whole visual field.

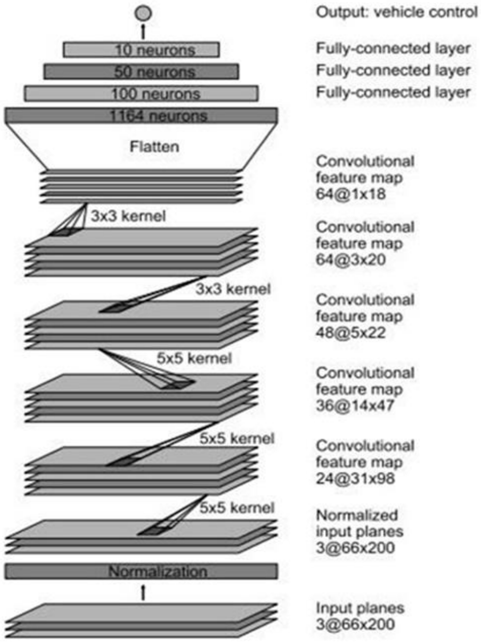

Fig. 4. ‘Pilotnet’ - Architecture Diagram

The model consists of the input layer (3D image vectors in the RGB space), a normalisation layer that makes the vector values more manageable, five convolutional layers that work with internal kernel filters which capture the salient regions of interest in the image and 3 dense layers that are used to help the model train via a method known as ‘Backpropagation’ (iterative gradient descent) followed by the output node.

B. Methodology

The figure below is a good illustration of the approach used in this project. As shown, the data (in the form of images) is collected and recorded in a folder with its corresponding log file (in the ‘.csv’ format). These files are then combined to produce the model (as described in the low-level design section). The model is saved in the standard ‘.h5’ format and reloaded later to make the predictions. The latter are passed into the simulator by way of opened socket connections and the car is driven based on these predictions.

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified ⋇ Impact Factor 7.39 ⋇ Vol. 11, Issue 7, July 2022
DOI: 10.17148/IJARCCE.2022.11789

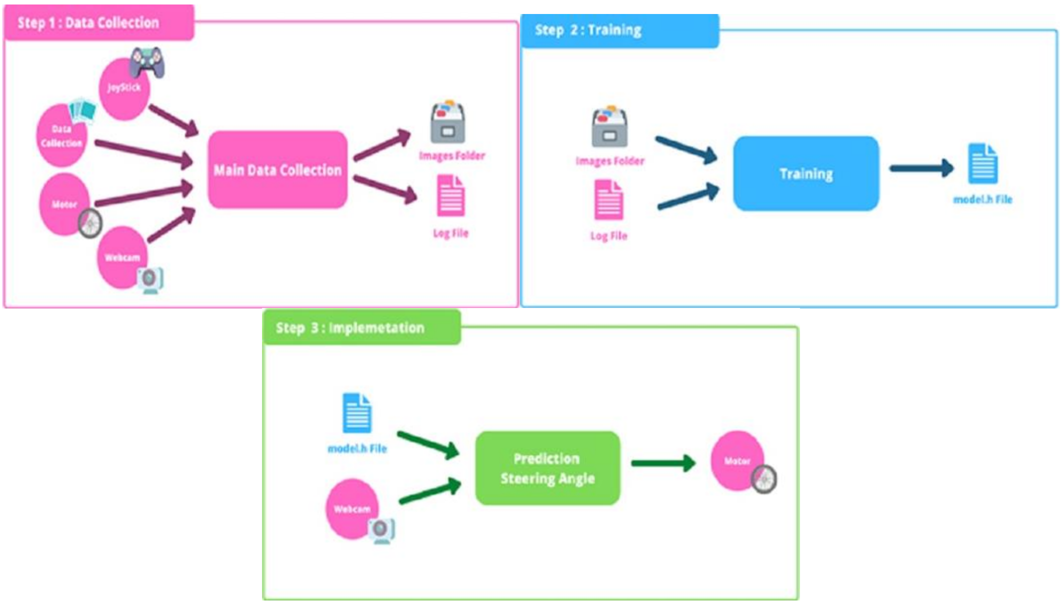

Fig. 5. Steps involved in training and deployment of the model in simulator

The architecture described, when deployed in track-1 of the Udacity simulator, generalizes quite well from the trained model on track-2. However, the model when deployed on track-2 performs quite poorly owing to the complex manoeuvres that have to be taken. In addition, the track is full of considerable inclines, declines and poor lighting.

The following is proposed to increase the accuracy of the final model:

1. Proportional Integral Derivative (or PID) controller for speed manipulation
2. Recovery data for specific badly performing manoeuvres
3. Transfer Learning using a pre-trained model from an older track with similar features (sharp changes in direction, inclines and declines)

PID stands for Proportional-Integral-Derivative in a PID controller. The best possible combination of these three elements results in the production of a control signal. The "proportional" component makes the vehicle react in a way that is opposite of and proportional to the inaccuracy in relation to the intended variable (in this case, speed). Instead of aiming to maintain a steady speed, the car therefore accelerates on an ascent and slows down on a downhill. When adjusting the speed, the "differential" component combats the P component's propensity to ring and overshoot. A bias in the P component is offset by the "integral" component so that the controller can reach the desired speed.

The incorporation of recovery data facilitated learning of much improved features that have built in safety mechanisms. The model learns to avoid the pitfalls of its previous iteration and drives more effectively around sharp bends.

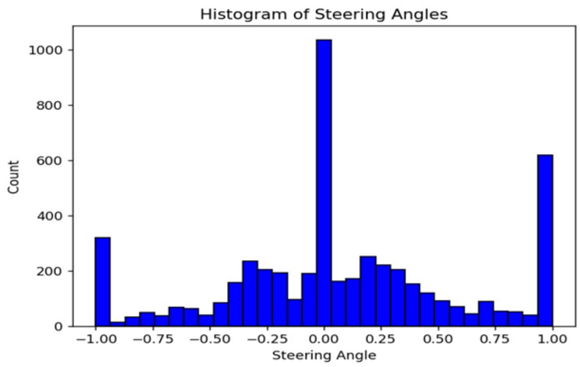

Fig. 6. Histogram depicting distribution of recovery steering angles in Track-2

DOI: 10.17148/IJARCCE.2022.11789

The basic premise of Transfer Learning is as follows: to utilize a different model with pre-trained weights for a similar but more general dataset and use it as a starting point for training. The model used for transfer learning is one that has been trained extensively on Track-3 shown in the previous section. The learned features are 'transferred' to the new model. Deploying this new model in the simulator shows considerable improvement over the earlier model as the vehicle is able to navigate the sharp bends, inclines and declines of Track-2.

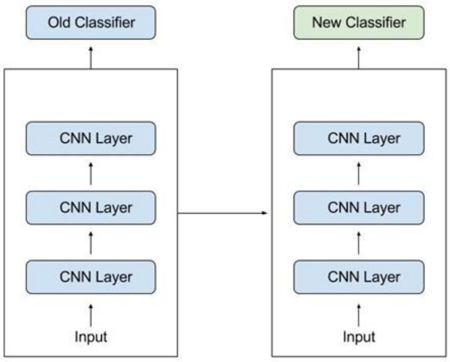

Fig. 7. Process of Transfer Learning

C. **Testing and simulation**

The problem domain in question is an example of a regression-based predictive domain. Here the objective is to predict numeric values like steering control and throttle and brake of the vehicle. It is different from classification-based predictive domain that involves predicting a class label. Therefore, unlike in classification, prediction accuracy is not a good or even relevant metric to measure the performance of the model. The skill or performance of a regression model is reported as an error in those predictions that measure how close they were to the expected values. Among error metrics, the 'mean squared error' (or MSE) is the most popular one. It is calculated as the mean or average of the squared differences between predicted and expected target values in a dataset. The system was tested by monitoring the training loss (predicted steering angles vs actual steering angles) and the validation loss (by running the trained model against the test set generated by the splitting process). The following plot was made for the loss tracked over the training and testing process:

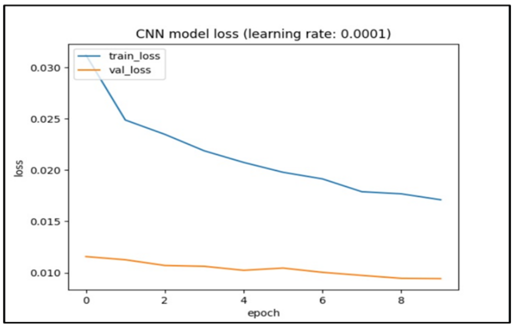

Fig. 8. Training and Validation (testing) loss – progression per epoch

V. Conclusion

It has been shown that Deep Learning is indeed a viable approach in developing autonomous driving units. This approach is simpler, while simultaneously providing good accuracy when trained with a reasonably sufficient amount of data. The incorporation of recovery data facilitated learning of much improved features that have built in safety mechanisms. The model learns to avoid the pitfalls of its previous iteration and drives more effectively around sharp bends.

VI. Results

The following graph illustrates the distribution of actual predictions made for a batch of input images and their corresponding expected values.

IJARCCE ISSN (O) 2278-1021, ISSN (P) 2319-5940

International Journal of Advanced Research in Computer and Communication Engineering

ISO 3297:2007 Certified ⁙ Impact Factor 7.39 ⁙ Vol. 11, Issue 7, July 2022

DOI: 10.17148/IJARCCE.2022.11789

As can be seen from the illustration, the distributions are similar albeit with minor variations.

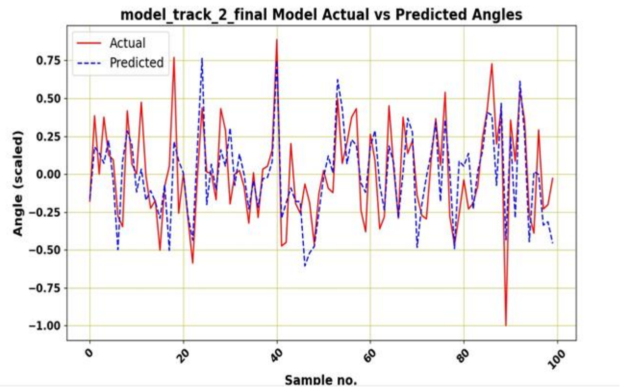

Fig. 9. Final model – Predicted values vs Expected values

The above posits Deep Learning as an effective tool in analysing input images for relevant features and making predictions regarding actuator controls on the basis of these learned features. However, the model is still quite simplistic to be directly deployed in actual roads. Testing must be carried out comprehensively to ensure that the vehicle understands recovery behaviour and proper road etiquette. Indeed, these deep learned features must be supplanted with further complexity such as planned routes and emergency safety mechanisms before the vehicle can be deployed in an active terrain. This Deep Learning approach can be contrasted with the 'Hierarchical approach' with many interacting modules as shown below. The latter is a complex model with immense overhead in terms of computation and feedback. Research is found to be more partial towards Deep Learning based architectures such as the 'Pilotnet' implemented in this project, because of their simplicity and understandability. Nevertheless, if Level-5 autonomy in driving is to be achieved, both approaches will have to be exhaustively studied and iteratively improved for the foreseeable future.

ACKNOWLEDGMENT

We are indebted to, **Prof. Merin Meleet,** Assistant Professor, Dept. of Information Science and Engineering, for her wholehearted support, suggestions and invaluable advice throughout our project work and also helped in the preparation of this thesis. Our sincere thanks to **Dr. Sagar B M**, Professor and Head, Department of Information Science and Engineering, RVCE for his support and encouragement.
We express sincere gratitude to our beloved Principal, **Dr. K. N. Subramanya** for his appreciation towards this project work.We thank all the teaching staff and technical staff of the Information Science and Engineering department, RVCE for their help. Lastly, we take this opportunity to thank our family members and friends who provided all the backup support throughout the project work.

REFERENCES

[1]. D. Pomerleau, "ALVINN: An Autonomous Land Vehicle in a Neural Network", Proceedings.neurips.cc, 1988. [Online]. Available: https://proceedings.neurips.cc/paper/1988/hash/812b4ba287f5ee0bc9d43bbf5bbe87fb-Abstract.html.
[2]. Bojarski, Mariusz, et al. "Explaining how a deep neural network trained with end-to-end learning steers a car." arXiv preprint arXiv:1704.07911 (2017).
[3]. Chen, Long, et al. "Deep neural network based vehicle and pedestrian detection for autonomous driving: a survey." IEEE Transactions on Intelligent Transportation Systems 22.6 (2021): 32343246.
[4]. Smolyakov, M. V., et al. "Self-driving car steering angle prediction based on deep neural network an example of CarND udacity simulator." 2018 IEEE 12th international conference on application of information and communication technologies (AICT). IEEE, 2018.
[5]. Do, Truong-Dong, et al. "Real-time self-driving car navigation using deep neural networks." 2018 4th International Conference on Green Technology and Sustainable Development (GTSD). IEEE, 2018.
[6]. Chernikova, Alesia, et al. "Are self-driving cars secure? evasion attacks against deep neural networks for steering angle prediction." 2019 IEEE Security and Privacy Workshops (SPW). IEEE, 2019.
[7]. Bechtel, Michael G., et al. "Deeppicar: A low-cost deep neural network-based autonomous car." 2018 IEEE 24th international conference on embedded and real-time computing systems and applications (RTCSA). IEEE, 2018.
[8]. Shengbo, L. I., et al. "Key technique of deep neural network and its applications in autonomous driving." Journal of Automotive Safety and Energy 10.2 (2019): 119.

International Journal of Advanced Research in Computer and Communication Engineering
ISO 3297:2007 Certified ⁜ Impact Factor 7.39 ⁜ Vol. 11, Issue 7, July 2022
DOI: 10.17148/IJARCCE.2022.11789

[9]. Haq, Fitash Ul, et al. "Comparing offline and online testing of deep neural networks: An autonomous car case study." 2020 IEEE 13th International Conference on Software Testing, Validation and Verification (ICST). IEEE, 2020.
[10]. Rao, Qing, and Jelena Frtunikj. "Deep learning for self-driving cars: Chances and challenges." Proceedings of the 1st International Workshop on Software Engineering for AI in Autonomous Systems. 2018.
[11]. Zaghari, Nayereh, et al. "Improving the learning of self-driving vehicles based on real driving behavior using deep neural network techniques." The Journal of Supercomputing 77.4 (2021):
[12]. Jalali, Seyed Mohammad Jafar, et al. "Optimal autonomous driving through deep imitation learning and neuroevolution." 2019 IEEE International Conference on Systems, Man and Cybernetics (SMC). IEEE, 2019.
[13]. Rausch, Viktor, et al. "Learning a deep neural net policy for end-to-end control of autonomous vehicles." 2017 American Control Conference (ACC). IEEE, 2017.
[14]. Chandrasekaran, Jaganmohan, et al. "A Combinatorial Approach to Testing Deep Neural Network-based Autonomous Driving Systems." 2021 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW). IEEE, 2021.
[15]. Kocić, Jelena, Nenad Jovičić, and Vujo Drndarević. "An end-to-end deep neural network for autonomous driving designed for embedded automotive platforms." Sensors 19.9 (2019): 2064.
[16]. Bojarski, Mariusz, et al. "End to end learning for self-driving cars." arXiv preprint arXiv:1604.07316 (2016).

