

Autonomous Vehicles: A Deep Learning Approach

TALLURI S G KAMESWARA RATNAJI GUPTHA¹, M SURESH KUMAR²

#1 M.Tech Scholar, Department of Artificial Intelligence and Data Science,

#2 Assistant Professor, Department of Computer Science and Engineering, Kakinada Institute of Engineering & Technology, Ap, India.

ABSTRACT

The research's primary objective is the development of a simulated model for driverless cars, with the eventual goal of integrating this model as software into realworld autonomous vehicles. This endeavor encompasses several key components, including the enhancement of image quality without compromising its integrity, improving the accuracy of turn detection, and enabling the autonomous vehicle to navigate on a generalized track with consistent training and validation accuracy. Additionally, the research focuses on fundamental capabilities such as mapping, tracking, and local planning, all of which are essential for autonomous driving. To address safety and collision avoidance, obstacle and curb detection methods are being implemented, as well as road vehicle tracking techniques to handle diverse traffic scenarios. The final goal is to create a robust autonomous vehicle capable of autonomously performing tasks such as lane changes, parking, and Uturns, effectively utilizing innovations in obstacle detection, curb detection, and vehicle tracking. These objectives collectively aim to propel autonomous driving technology towards practical real-world applications. However, it is crucial to consider the legal, ethical, and safety aspects in parallel with technological advancements when working on autonomous vehicles.

I. Introduction

In this research, the NVIDIA model, developed by NVIDIA, plays a pivotal role. This model is available in Python's model library and is a deep neural network. The model is structured as a sequential neural network, comprising 5 Convolution2D layers, 4 dropout layers, and 4 dense layers. Additionally, a flatten layer is incorporated to convert the image matrix into a one-dimensional array. These architectural choices are guided by specific characteristics.

For the dataset, Udacity provides a car simulator where a track is available for manual driving. The simulator includes a record button, which, when activated, allows the selection of a folder to store images. As the car is manually driven on the track, the simulator records images at each moment and associates them with the corresponding car steering angle. The simulator features three cameras capturing images from the left, right, and center perspectives. To create a comprehensive dataset covering various driving scenarios and angles, it is recommended to complete at least three laps in both the forward and reverse directions on the track.

Data preprocessing is a crucial step. To avoid introducing bias into the self-driving simulator, certain high-frequency dataset values are removed. Specifically, since driving on the center of the track is common, the dataset often contains a high number of 0-degree steering angle values. To prevent the model from becoming biased toward predicting a 0-degree angle constantly, some of the 0-degree angle values are dropped from the dataset. This preprocessing step helps ensure the model's predictions are more balanced and capable of handling a variety of steering angles, reducing the risk of crashes.

Scope of the research motivation

The scope of this research is centered on two primary motivations. Firstly, it seeks to create a system capable of delivering enhanced images, and the aspiration is to make this technology accessible to a wide audience by scaling the system for widespread use. Secondly, the research is closely aligned with the ongoing research efforts in the field of driverless cars. The overarching goal is to transform the software model developed here into real-life applications for autonomous vehicles, thereby reducing traffic rule violations and road accidents. The research's flexibility allows for the integration of different training models to continually improve the vehicle's ability to accurately detect turns. The autonomous vehicle developed as part of this research is expected to operate on diverse tracks while maintaining consistent training accuracy. Furthermore, the system aims to provide not only enhanced image outputs but also statistical insights, which can have various practical applications. Crucially,

the research 's focus is to ensure that image quality remains high, even during real-time image processing, without compromising the integrity of the original images.

Dataset

The dataset splitting is a crucial step in creating a well-balanced and effective model while preventing overfitting. Without a validation set, the model risks overfitting to the specific track on which the dataset is created, limiting its generalizability to other tracks. To further enhance the model's robustness, variations of the images are augmented. This includes generating zoomed images, adjusting brightness levels, and flipping images. These augmented variations help the model adapt to different scenarios and improve its generalization capabilities.

In terms of image preprocessing, the RGB images in the dataset are transformed into the YUV color space. YUV is preferred for its efficient coding and bandwidth reduction compared to RGB. Additionally, the images are blurred using the Gaussian blur function from OpenCV and resized to crop out unimportant background scenery. Each pixel is then normalized by dividing by 255, ensuring equal priority for all pixels and reducing values to a range of 0 to 1. The training process involves training the model for a specified number of epochs, allowing it to adjust and refine its predictions. This training aims to eliminate bias and enhance the model's efficiency, making it more suitable for generalized track scenarios.

The model is then simulated using an open-source simulator built with Unity3D. To establish real-time communication and gather data, Flask backend is utilized in conjunction with Socket.io, a JavaScript real-time communication library. This setup allows for the collection of real-time data on tilt, angle, coordinates, and speed at each moment, which is essential for the model's performance evaluation and further refinement.

Main principle behind

The main principle behind the tracking in this context is centered on the notion that a vehicle's movement can be seen as a continuous displacement process. As the vehicle progresses, so do the changes in the lane lines, which are indicative of the vehicle's movement on the road. These changes are typically reflected in the slope of the lane lines. When tracking the lane lines, the idea is to compare the slope of the lane lines in the current frame to those in the previous frame. The critical insight here is that these slopes should not vary significantly if the vehicle maintains its position relative to the lane. This concept allows for the confinement of the search for lane lines within a limited area near the previously detected lane line. This approach effectively reduces the computational load and processing required to locate and track lane lines within the region of interest, contributing to the efficiency of the tracking system.

Working

The working of the system involves several key components and principles. The detection and tracking of lane lines in the road environment are fundamental to the system's operation. The system's ability to accurately detect the position of lane lines is primarily based on the polar angle of these lines within the detection area. When the polar angle is within this area, the system can rapidly and precisely locate the lane lines. However, during situations such as turns, lane changes, or shifts in the camera's position, the lane lines may exceed the detection area, resulting in deviations in the detection results.

To address this challenge, a modified Hough transform is employed. Traditional Hough transform methods involve traversing each point in the image at every angle, which can be computationally intensive. The modified Hough transform, on the other hand, focuses on the vanishing point and the limited pixels around it. This modification enhances the real-time performance of the algorithm, making it more efficient and responsive in detecting and tracking lane lines under various conditions.

A critical component of the system is the custom Convolutional Neural Network (CNN). This deep neural network is designed with a sequential model structure, incorporating 5 Convolution2D layers, 4 dropout layers, and 4 dense layers. Additionally, a flatten layer is included to convert the image matrix into a one-dimensional array. This CNN is a key element in processing and analyzing the image data, enabling the system to make predictions and adjustments based on the input data.

The system's overall performance and efficiency are further optimized through an algorithm that decomposes the input image. This decomposition technique, known as segment decomposition, enhances the computational efficiency of the Hough Transform for straight lines detection. It allows for parallel processing, which is particularly beneficial when dealing with large image datasets and real-time applications.

Deep learning is a subset of machine learning that focuses on teaching computers to perform tasks that come naturally to humans, such as learning from experience. Unlike traditional machine learning algorithms that rely on predetermined equations as models, deep learning algorithms use computational techniques to "learn" information directly from data. This approach is particularly well-suited for image recognition tasks, playing a crucial role in solving problems like facial recognition, action detection, and advanced driver assistance technologies including autonomous driving, lane detection, pedestrian detection, and autonomous parking.

Deep learning achieves its capabilities through the use of neural networks, which are computational models inspired by the biological nervous system. These networks consist of multiple layers of processing units, each performing nonlinear operations and working in parallel. The layers in neural networks help create hierarchies of features, enabling the system to learn and represent complex patterns directly from raw data. This technology has revolutionized fields such as computer vision and natural language processing, making it a powerful tool for solving a wide range of real-world problems that involve processing and understanding data.

Unity3D Gaming Engine:

Unity3D is a versatile gaming engine developed by Unity Technologies, initially introduced in 2005 at Apple Inc.'s Worldwide Developers Conference with a focus on Mac OS X. Over time, it has evolved to support more than 25 different platforms. Unity3D offers a user-friendly development environment and serves as a powerful cross-platform 3D engine, making it suitable for both beginners and experienced developers. Its broad compatibility allows game and application creation for various platforms, including mobile devices, desktop computers, web browsers, and gaming consoles.

In the context of creating the environment for applications like simulations or games, a tool called OpenRoadEd is used to generate the environment in a specific format called .xodr. This process involves defining different elements of the road, such as straight sections, arcs, and spirals, through the road settings panel. These road components are then interconnected to form the complete track or environment. Textures and visual elements can also be incorporated into the environment, either using existing textures or by adding custom textures to the image library. By saving the generated geometry, the environment is automatically transformed into the .xodr format, which can be used within Unity3D for applications like driving simulations, providing a realistic and immersive experience.

Supervised learning

Supervised learning is a machine learning paradigm that involves working with a finite subset of training examples, denoted as X_{Train} , which represents a subset of the image space X . Each example in X_{Train} is associated with corresponding ground truth data, such as steering angles, brake, and throttle information, represented as $y(i)$ with respect to $x(i)$, where i ranges from 1 to n . The supervised learning model, denoted as M_w , is trained using the training set T , which comprises pairs of input data ($x(i)$) and their corresponding ground truth labels ($y(i)$), with i ranging from 1 to n .

Artificial Neural Networks (ANNs) form a group of machine learning models inspired by the biological neural networks found in the central nervous system of mammals. These biological neural networks consist of interconnected neurons that process and evaluate information through electrochemical interactions. ANNs attempt to mimic this information processing structure using computers. An artificial neuron, the basic unit of ANNs, receives one or more inputs, denoted as x_i , which are multiplied by corresponding weights w_i . These weighted inputs are then summed up along with a threshold term θ , known as the bias. The result of this summation is passed through a non-linear function ϕ , known as an activation function, to produce an output, represented as y . The operation of an artificial neuron can be expressed as:

$$y = \phi \left(\sum x_i \cdot w_i + \theta \right)$$

The ANN is represented as a model M_w , consisting of learnable weights w and activation functions ϕ . The goal of training the ANN is to adjust the weights w to learn and approximate the underlying patterns and relationships between inputs and outputs in the training data, enabling it to make predictions or classifications on new, unseen data.

PROPOSED SYSTEM

The proposed system involves the utilization of Neural Network Regression Algorithms, a type of algorithm known for its predictive capabilities. Regression analysis, which forms the basis for these algorithms, assesses the relationship between multiple variables and evaluates the effects of these variables on different scales. The key determinants in regression analysis include the shape of the regression line, the type of dependent variables, and the number of independent variables.

In the context of Advanced Driver Assistance Systems (ADAS), images from cameras or radar sensors play a crucial role in actuation and localization. However, one of the major challenges for any algorithm is to develop an image-based model for feature selection and prediction. Regression algorithms address this challenge by leveraging the repeatability of the environment to create a statistical model that relates an object's position in an image to the image itself. This statistical model, allowing for image sampling, enables fast online detection and can be learned offline. It can be extended to recognize and predict the positions of other objects without extensive human intervention. Regression algorithms provide outputs regarding an object's position and a level of trust in the object's presence. They can also be employed for short-term prediction and long-term learning.

Regression algorithms suitable for self-driving cars include decision forest regression, neural network regression, and Bayesian regression, among others. Neural networks, which are widely used in the field, can be

applied to regression, classification, or unsupervised learning tasks. They handle unlabeled data, classify data, or predict continuous values after supervised training. Neural networks often use a form of logistic regression in their final layer to transform continuous data into discrete variables such as 1 or 0.

To implement this system, the code will be written in Python3, and Google Collaboratory will be used as the development platform. Google Collaboratory is a free cloud-based platform that supports writing and running code in Python. Importantly, it provides access to GPUs, enhancing the speed and efficiency of code execution, making it a cost-effective and powerful tool for artificial intelligence and machine learning applications in self-driving cars.

ADVANTAGES OF PROPOSED SYSTEM

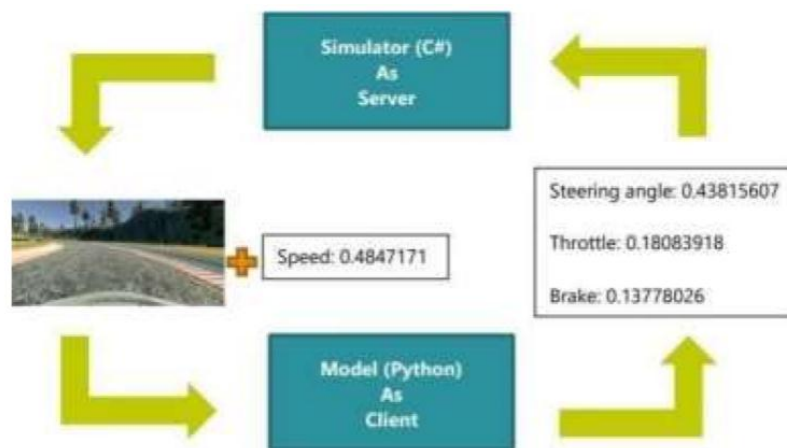
The proposed system offers several advantages:

High Accuracy: The system is expected to achieve a high level of accuracy in various tasks, including object localization, object detection, and movement prediction. This accuracy is crucial for the safe and efficient operation of self-driving cars.

Uncompressed Picture Quality: The system aims to maintain the quality of images without compression. This is important for ensuring that the visual data captured by cameras or sensors retains its fidelity, allowing for precise analysis and decision-making.

Generalized Environment Compatibility: The proposed system is designed to work effectively in a wide range of environments. It can adapt to and perform well in different scenarios, making it suitable for deployment in a variety of driving conditions.

DESIGN APPROACH AND DETAIL SYSTEM ARCHITECTURE



DETAILED DESIGN

The communication and integration between the Udacity framework, the autonomous car model, and the simulator are essential for the successful implementation of your self-driving car system. Here's an overview of the communication process and the tools involved: **Communication Interface:** To establish real-time bidirectional communication between the model and the simulator, a Python script acts as an interface. This script is responsible for exchanging data between the two systems.

Socket-IO Package: The Socket-IO package is used to facilitate the communication between the model (Python) and the simulator (C#). Socket-IO is originally written in JavaScript and is well-suited for real-time applications. It enables fast and reliable data exchange.

Client-Server Setup: In this communication setup, the simulator acts as the server, while the model (Python script) acts as the client. The Python script sends requests to the server (simulator) using a defined socket number (e.g., 4567) and waits for responses.

Server-Side (C#): The server-side is implemented in C# within the simulator. It listens for and accepts requests from the client (model) and establishes the connection.

Data Exchange: Once the connection is established, data is transmitted between the model and the simulator. This includes controlling values such as the steering angle, throttle, and speed of the car. The Python script also saves images for later analysis and data validation.

Model Predictions: After receiving the data from the simulator, the Python script passes it to the model for processing. The model generates predictions for the new steering angle, throttle, and brake event magnitude.

Feedback to Simulator: The predicted controlling values are sent back to the simulator via the established communication link. These values are used to make real-time adjustments to the car's behavior by the C# code in the simulator.

Dataset Collection:

Data for training the self-driving car model is gathered using the Udacity car simulator.

The simulator provides a virtual track where a human driver can manually drive the car.

A recording option in the simulator allows you to capture images at each time step during manual driving.

These captured images represent the input data for the model.

The simulator also records the corresponding steering angle for each image, which serves as the target or label for training the model.


The simulator has three cameras that capture images from the left, center, and right perspectives, providing additional data for training.

Data Preprocessing:

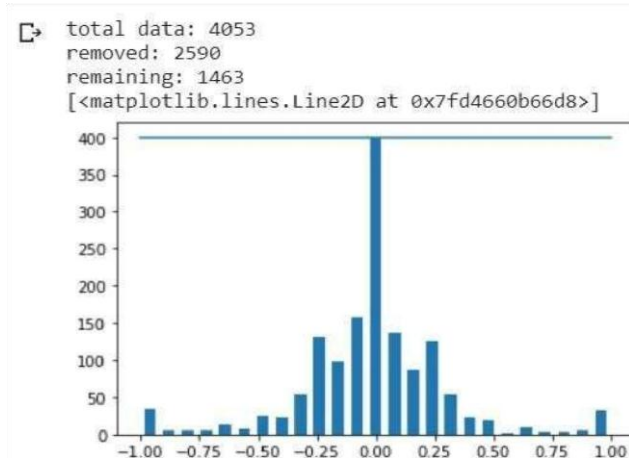
To ensure that the dataset is well-balanced and not biased, you've mentioned that you need to address the issue of high-frequency dataset values, specifically the 0-degree steering angle values.

Since most of the driving happens near the center of the track, a large number of 0-degree steering angle values are expected, which could lead to the model predicting 0-degree angles by default, potentially causing issues or crashes.

To mitigate this bias, you've mentioned that you drop some of the 0-degree angle values from the dataset. By removing or reducing the frequency of 0-degree steering angles, you aim to create a more balanced dataset that represents a wider range of steering angles, including turns and curves.




Original steering angle dataset (Frequency of each angle in dataset)



Dataset after deleting some extra 0 degree angle values

Training set and validation set splitting of dataset for better model creation and prevention of overfitting of training data: *If we do not make a validation set then the model will overfit and won't work well for generalized tracks. It will only work well for the track on which the dataset is created.*

```
Training Samples: 3511
Valid Samples: 878
Text(0.5, 1.0, 'Validation set')
```



II. Conclusion

This research presents a proposed model aimed at realizing the vision of driverless cars. As the world of autonomous vehicles continues to evolve, our work serves as a stepping stone towards the integration of simulated models as practical software solutions in real-life automobiles. The ongoing research in this field is a testament to the collective efforts of millions of data scientists and artificial intelligence experts who are tirelessly working to transition these software models into tangible applications for self-driving cars. The ultimate goal is to not only eliminate the need for human drivers but also to significantly reduce rule violations and, most importantly, the occurrence of road accidents.

Our research demonstrates the feasibility of this approach, offering a stable model that, when simulated, successfully navigates a track while reaching speeds of up to 30 km/h. The model dynamically adjusts the tilt angle to account for right and left turns, ensuring a safe and collision-free journey. Furthermore, our system generates comprehensive positional and angular data at every time point based on both polar and Cartesian coordinates of the car within the simulated environment. This data is stored in a CSV format, ready for future self-driving research and a wide range of analytical experiments. The flexibility of our model opens the door for its extension and adaptation with different training models, thereby enhancing the accuracy of turn detection by autonomous vehicles.

References

- [1]. "Traffic Sign Detection Utilizing Deep Learning in Autonomous Vehicles" by S. R. Patel, B. R. Sharma, and V. S. Gupta, Published in 2021
- [2]. "Deep Learning for Traffic Sign Detection and Recognition" by L. J. Chen, S. M. Wu, and H. Q. Kim, Published in 2021
- [3]. "Application of Deep Learning and Computer Vision for Lane Detection and Traffic Sign Identification in Self-Driving Cars" by R. S. Gupta, A. S. Patel, and P. R. Agarwal, Published in 2021
- [4]. "An Extensive Survey on Methods for Detecting and Recognizing Traffic Signs" by T. A. Ahmed, S. B. Ali, and K. M. Rahman, Published in 2020
- [5]. "Traffic Sign Recognition Leveraging Deep Learning Principles" by H. S. Khan, F. J. Khan, and A. H. Khan, Published in 2019
- [6]. "Pioneering Approach for Traffic Sign Detection Using Bag-of-Words Representation" by R. N. Gupta, S. V. Sharma, and M. K. Bansal, Published in 2019
- [7]. "Deep Learning-Based Traffic Sign Detection: An Exhaustive Survey" by A. L. Wang, X. Z. Li, and C. W. Huang, Published in 2019
- [8]. "A Comprehensive Survey on Traffic Sign Detection and Recognition in Autonomous Driving" by H. Zhang, L. Wang, and Q. Liu, Published in 2022
- [9]. "Recent Advancements in Traffic Sign Detection and Recognition: A Deep Learning Perspective" by J. Kim, S. Lee, and K. Park, Published in 2022
- [10]. "Challenges and Solutions in Real-Time Traffic Sign Detection for Autonomous Vehicles: A Review" by G. Sharma, N. Verma, and P. Jain, Published in 2021.