

Review

Object detection, distributed cloud computing and parallelization techniques for autonomous driving systems

Edgar Cortés Gallardo Medina¹, Victor Miguel Velazquez Espitia¹, Daniela Chípuli Silva¹, Sebastián Fernández Ruiz de las Cuevas^{1,2}, Marco Palacios Hirata¹, Alfredo Zhu Chen¹, José Ángel González González¹, Rogelio Bustamante Bello¹ and Carlos Francisco Moreno-García^{3,*}

¹ Tecnológico de Monterrey, School of Engineering and Sciences. Calle Puente 222, Coapa, Arboledas del Sur, Tlalpan, 14380, Mexico City, Mexico.

² Politecnico Di Torino, Department of Mechanical and Aerospace Engineering. Corso Duca degli Abruzzi, 24 10129 Torino, Italy

³ Robert Gordon University, School of Computing. Garthdee Road, AB10 7QB, Aberdeen, Scotland, UK.

* Correspondence: c.moreno-garcia@rgu.ac.uk

Received: date; Accepted: date; Published: date

Featured Application: This paper contains a comprehensive review of the current state of the art machine vision algorithms used to build a pipeline towards achieving assisted driving and fully autonomous vehicle development. Key and novel concepts such as Distributed Systems (DS) and parallelization are explored to enhance said systems' functionality.

Abstract: Autonomous driving systems are increasingly becoming a necessary trend towards building smart cities of the future. Numerous proposals have been presented in recent years to tackle particular aspects of the working pipeline towards creating a functional end-to-end system, such as object detection, tracking, path planning, sentiment or intent detection. Nevertheless, few efforts have been made to systematically compile all of these systems into a single proposal that effectively considers the real challenges these systems will have on the road, such as real-time computation, hardware capabilities, etc. This paper has reviewed various techniques towards proposing our own end-to-end autonomous vehicle system, considering the latest state on the art on computer vision, DSs, path planning, and parallelization.

Keywords: Autonomous driving systems, computer vision, neural networks, feature extraction, segmentation, assisted driving, cloud computing, parallelization.

1. Introduction

The National Highway Traffic Safety Administration (NHTSA) reported that 94% of severe crashes in the road are caused by human error [1]. In this regard, the rise of the autonomous vehicle (AV) has a huge potential to decrease these accidents and make the road much safer. Therefore, the implementation of robust and secure systems is paramount for the proper design of an autonomous driving system (ADS) pipeline. This field has been widely investigated for many years. Both academia and industry have investigated that has achieved breakthroughs and state of the art results in the last years. In 2004, the Defense Advanced Research Projects Agency (DARPA) held the Grand Challenge competition for American AVs¹; however, no team could complete the designated 150 miles route. One year later at the same event, all but one of the finalists surpassed the best record

¹ DARPA, *The Grand Challenge*, DARPA, <https://www.darpa.mil/about-us/timeline/-grand-challenge-for-autonomous-vehicles>

(7.32 miles) from the previous year, and five teams were able to finish the 132 miles desert terrain route without human intervention. In 2007, during the third edition of this challenge², 60 miles of the urban course were proposed. Only six teams managed to complete the entire track successfully. These three competitions were very significant events up to this time, so the results have inspired universities and large corporations to improve the state of the art of in ADSs in different environments and conditions.

Apart from the DARPA challenges, many other remarkable competitions have been held until today. For instance, Hyundai Motor Company has been researching AVs for a long time. They aim to expand future vehicle research at universities and improve Korea's automotive industry's technological development, and thus they have organized four AVs competitions³ since 2010. Similarly, the Society of Autonomous Engineers (SAE) and General Motors have partnered to sponsor the AutoDrive Challenge⁴, which consisted of a three-year (2018-2020) competition for students to complete an urban driving course. Furthermore, Indy Autonomous Challenge⁵ is a new high-speed autonomous race held at the Indianapolis Motor Speedway. Every team must use the specific Dallara-produced IL-15 that has been retrofitted with hardware and controls to compete. It has been noted that numerous universities and companies worldwide are very active in these kinds of events, which indubitably will accelerate research on this field.

From the academic perspective, several survey papers in this area have taken different approaches to review this field. Rangesh et al. [2] compared several representative online Multi-Object Tracking (MOT) for AVs. Later on, authors proposed M³OT (i.e. a 3D variation of MOT), an approach suitable for tracking objects in the real world. It is worth mentioning that M³OT can work with any type and number of cameras as long as they are calibrated. Huang et al. [3] spanned the state of the art technology in significant fields of self-driving systems, such as perception, mapping and localization, prediction, planning and control, simulation, etc. Yurtsever et al. [4] attempts to provide a structured and comprehensive overview of state of the art automated driving-related hardware-software practices, high-level systems architectures, present challenges, datasets and tools related to ADS.

Badue et al. [5] published a survey research on self-driving cars based on the aforementioned DARPA challenges. The authors presented a detailed description of the self-driving vehicle developed at the Universidade Federal do Espírito Santo in Brazil, named Intelligent Autonomous Robotics Automobile (IARA). This paper presents a complete overview of each of the aspects that composes the matured used architecture of self-driving cars nowadays, dividing it into two systems, i.e. the perception and decision-making. At the same time, each design is composed of a variety of subsystems. Achieving its objective by describing each of the technologies superficially, this work represents a notable introduction to the interested reader. However, for a lettered person in the subject, this work's scope might be short, because it doesn't present the state of the art technologies for the different paradigms of driverless cars.

² DARPA, *DARPA Urban Challenge*, DARPA, <https://www.darpa.mil/about-us/timeline/darpa-urban-challenge>

³ Hyundai NGV Industry-Academy-Research Cooperation, *Autonomous Vehicle Competition*, Hyundai NGV, http://www.hyundai-ngv.com/en/tech_coop/sub05.do

⁴ SAE International and GM, *AutoDrive Challenge*, SAE, <https://www.sae.org/attend/student-events/autodrive-challenge/>

⁵ Indianapolis Motor Speedway, *Indy Autonomous Challenge*, Indianapolis Motor Speedway, <https://www.indyautonomouschallenge.com/>

This paper provides a state of the art overview investigation on object detection, distributed cloud computing, and parallelization techniques for assisting driving and AVs. Unlike other reviews, we centralized our study mainly in autonomous driving perception systems, focusing on the concepts of the neural network (NN) for object detection and the DS for parallelization. The paper is structured as follows. Section 2 provides the basic definitions and background required to understand the domain, with concepts such as image acquisition, feature extraction, object detection, DS and route calculation. Section 3 presents a thorough study of each of these concepts in light of the latest work presented in this regard. In Section 4, we offer our proposed methodology to address the most recurrent issues of assisted driving and AVs based on our study and with considerable focus on parallelization for real-time deployment systems. Finally, in Section 5, we provide our final thoughts and comments on the subject.

2. Background

In this section, the object detection task's building blocks are explained along with their corresponding evaluation metrics and applications (i.e. path planning using the Frenét framework). We start by explaining feature extraction; the types and methods used to discern distinctive characteristics from the data provided by various sensors (also defined in this section). Moreover, we will cover the discussion of multiple NN architectures and the concatenation of several deep learning techniques, which will be used for the Frenét framework generation for path tracking of vehicles on the road. DS architectures, whose components are located in different networks, are also explored in deep learning systems. These can provide an advantage when training and running the computer vision architectures.

2.1. Feature Extraction

An object detection system's performance is determined not only by the learning algorithm responsible for most of the task at hand but also by the methods that select and combine variables into discernible features or distinctive attributes from images or other sensor data. In the end, the learning algorithm is as good as the quality of the features used to train it⁶. The raw sensor data itself contains structural information which can be extracted by the learning algorithm itself (such as in the case of the Convolutional Neural Network). However, there is usually a process undertaken to identify, extract and assign features so that the model learns only the most important characteristics and obtains better results. It is also important to select features based on the application and the relevance of the pertinent problem. This work discusses the most common feature extraction approaches in machine learning for multiple object detection, such as camera-based, Light Detection and Ranging (LiDAR) based, and a hybrid method that uses both. Although NN architectures can be used for feature extraction as wellpoint-cloud method level of abstraction, which can be useful depending on the application.

2.1.1. Camera-based

Most common commercial cameras can capture rich color and texture information from images used to locate objects and obstacles along the course. An image typically has a representation of three channels. On the one hand, attempting only to use these components' values is defined as a *pixel feature approach*; this means that the machine learning algorithm will obtain and learn the raw data features by itself. On the other hand, crafted feature representation consisting of image processing techniques to extract the features before the learning process, such as edge detection, corner detection,

⁶Vickery, R., The Art of Finding the Best Features for Machine Learning, *Towards Data Science*, [The Art of Finding the Best Features for Machine Learning | by Rebecca Vickery | Towards Data Science](#)

transformations, smoothing or blurring (among others), can be extracted before the learning process. This process is known as an *extracted feature approach*. The latter follows a series of traditional image pre-processing methods to extract a series of key points and its respective feature descriptors. For instance, Scale Invariant Feature Transform (SIFT) [6] is one of the most widely used methods for such purpose, because of its outstanding capabilities of extracting pixel colors, intensities, position, and Gaussian properties such as the Difference of Gaussian (DOG) pyramid and edge key points. All of this information is extracted while preserving scale invariance, as its name suggests. Although this is a suitable feature extractor for a camera's raw images, it is very slow in implementation and unusable in a real-time scenario. In contrast, the Speeded-Up Robust Features (SURF) [7] method is one of the fastest and most robust feature extractors. It behaves similar to SIFT, but its outstanding accuracy in extracting important key points between matching images, its real-time performance capability and the possibility for parallelization with GPU architectures using CUDA C/C++ libraries, puts this feature extractor at the forefront [8]. SURF uses Hessian matrices to detect blobs and to find important key points. Then, these are used to measure the focal length with each neighboring point in the image. Afterwards, the sum of Haar cascades is used to obtain rotational and translational invariance. Finally, its descriptors are compared to those on another image picturing the same object, but from a different angle.

2.1.2. LiDAR-based

LiDAR is one of the most widely used sensors in robotics [9]. The classical 2D LiDAR uses a rotating mirror, which sends laser beams along the x-y plane, sensing objects within range. On the other hand, a 3D LiDAR also shoots beams along the x-z plane, capturing the distance to items along this plane.

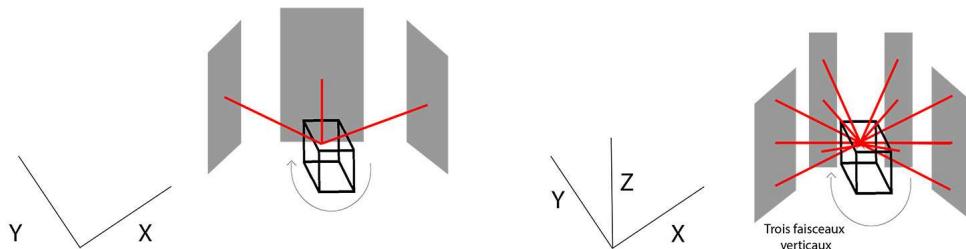


Figure 1. Diagram of the differences between the data gathered from 2D and 3D LiDAR [10].

The features extracted from a LiDAR sensor are the so-called *point clouds*. Point clouds represent the surfaces in which the LiDAR laser bounced before reaching the LiDAR sensor back. These points follow a specific array of algorithms to be scale, position and orientation invariant; all of this to perform Simultaneous Localization And Mapping (SLAM). LiDAR sensors can make a single point cloud swift in an outstanding short amount of time, which means that feature extraction is performed with a comparable speed to classical feature extraction. When it comes to object detection, a LiDAR sensor can also be used to aid AVs. BirdNet [11], for example, can detect 3D obstacles from the LiDAR's point cloud using three steps: first, a Bird's Eye View (BEV) image is created from the raw sensor data. Then, using CNNs intended initially for image processing, the position and orientation of obstacles are estimated. Finally, for the post-processing phase, predictions are corrected using statistical data, and the height of the barriers is projected on a predicted plane. This method was tested using the KITTI data set to show state of the art performance. Similar approaches include the Multi-View 3D Object Detection Network for Autonomous Driving (MV3D) [12]. MV3D fuses the camera and the LiDAR features to obtain a more rigorous representation, although the results of both methods show state of the art performance. When it comes to LiDARs in AVs, 3D LiDARs [9] are used for various reasons. First, this is a more robust sensor that includes more laser scans along the

z-axis so a more extensive range of obstacles can be detected, whereas a 2D LiDAR is bounded to scanning a fixed-line along this axis. Second, the 3D point cloud produced by a 3D LiDAR can easily be transformed to the camera's coordinate system. Thus, the object detection task can be performed by the camera and aided by the LiDAR as explored in the following section.

2.1.3. Hybrid

Several approaches intend to combine multiple modalities of features to provide complementary information and leverage the classification performance. The single modality of the features is not enough to solve some scenarios, such as self-driving, which involves volatile strategies and requires the utmost precision for safety reasons. In these cases, extracting features from different sensors and fusing them might improve obstacle detection effectiveness.

When it comes to self-driving cars [4], one common approach is to use the data of both cameras and LiDARs. Combining the cameras' detail to achieve reliable object detection with the laser scan reading provided LiDAR allows us to identify obstacles and place them in the world. As discussed in Section 2.1.2, a 3D LiDAR enables the fusion of a 3D point cloud and a 2D image, combining image detail and object positioning.

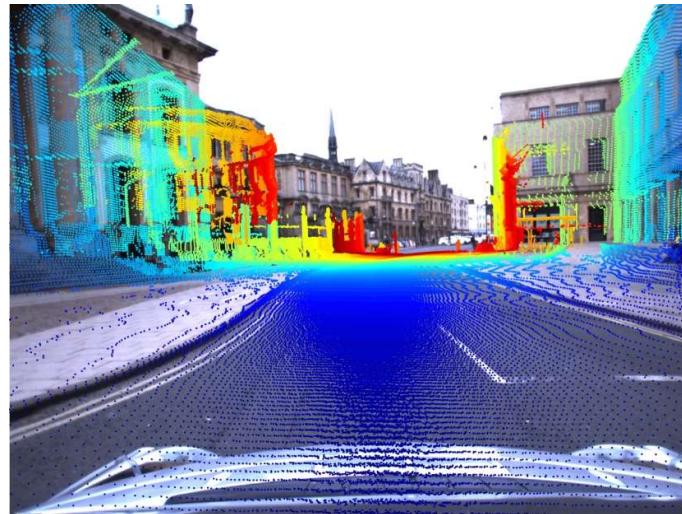


Figure 2. LiDAR point cloud projected on a 2D image [13].

To validate this, we ran a test inside a gazebo simulator, where the point cloud from a 2D LiDAR (mounted on a 2-wheel robot) was projected onto an image produced by a camera on the same robot. We projected the data from one sensor onto another, thus proving this method to be a low computation and low-cost procedure to accurately detect the distance to objects identified by computer vision algorithms. However, for the robustness needed in ADSs, the 3D LiDAR is still the more reliable option as the 2D LiDAR is bounded to a single line along the x-y plane.

It is essential to point out that using a LiDAR sensor is not the only ways to obtain point clouds. Using a stereo camera and performing SIFT-like algorithms, point clouds can also be generated for autonomous driving applications. The only issue is the frame rate, which is not as good as when using a LiDAR swift. Moreover, this alternative would require image pre-processing to acquire the point clouds properly. Nonetheless, this alternative is much less expensive, and even some GPU-based architectures of this kind to have shown a very competitive point cloud generation compared to LiDAR sensor-based architectures [14].

2.2. Convolutional Neural Network (CNN)

CNN is one of the most widely used image classification methods, object detection and semantic image segmentation because of its robust feature learning and classification capabilities [15], [16]. Their primary purpose is to reduce the images into a more straightforward process without losing features relevant for the right prediction. In other words, the objective is to extract the high-level features from the input image.

A typical CNN consists of an input layer, a convolutional layer, a pooling layer, a fully connected layer, and an output layer. The convolutional layer contains feature detectors, best known as kernels or filters, followed by a normalization layer and a non-linear activation function to extract feature maps of the input image. These feature detectors swift through the information in the input image and filter the critical parts to scan images and categorize them more accurately. These feature maps are then down-sampled by pooling layers, which remove redundant features and reduce their dimensions from its previous layers to improve the statistical efficiency and model generalization. Finally, fully connected layers are applied to classify the extracted features and obtain an input-based probability distribution. After that, the network's output is a fixed-sized vector [16], [17].

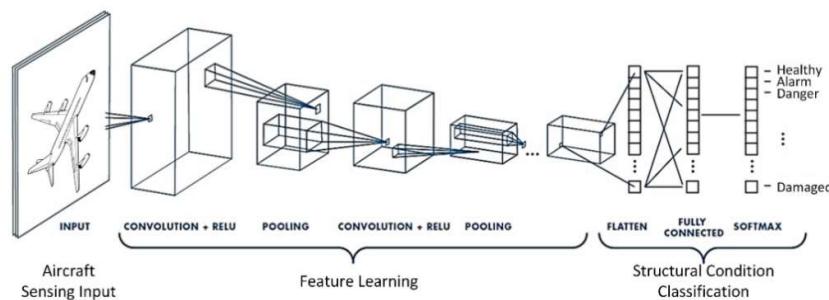


Figure 3. CNN architecture with an input, convolutional, pooling, fully connected, and output layers [18].

In terms of accuracy, the CNN is superior to traditional image classification methods, such as Directed Acyclic Graph-Support Vector Machine (DAG-SVM) [19] or Constrained Linear Discriminant Analysis (CLDA) [20], and are the most developed networks for image classification tasks nowadays [16], [17]. The structure of early CNNs was relatively simple, such as the classic LeNet-5 model [21] (shown in Fig. 4), which was mainly used in handwritten character recognition and image classification. However, with the continuous deepening of research and the need for more complex and extended applications, CNN's structure is continuously optimized. For example, the Convolutional Deep Belief Network (CDBN) [22] is an unsupervised generative model resulting from combining a CNN and a Deep Belief Network (DBN). Successful applications in the field of object detection (discussed in detail in the following section) were presented in [23] with the intent of performing region feature extraction based on a CNN.

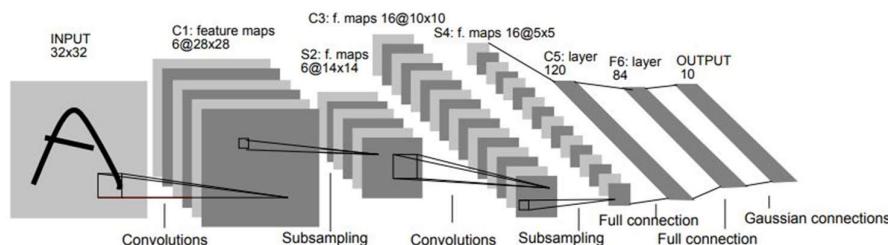


Figure 4. LeNet-5 architecture, which contains 5 first-level features [24].

2.3 Object Detection

Object detection is the action of recognizing and locating every single object within an image. In contrast, image classification detects all items in an image but only displays those with the highest probability of being the one of interest. Therefore, object classification is commonly used for classification tasks, while object detection is demanded in AV scenarios.

2.3.1 R-CNN

The Region-Based Convolutional Neural Network (R-CNN) model is an object detection architecture. This model allows selective search, which is based on proposing $\approx 2'000$ regions per image, which are then passed through the convolutional layers for the classification of the regions and features differentiation, to finally output the boundary box and the label of the object's type. [25] The principal disadvantage of this method is that it is very slow, taking around 50 seconds for object detection. Another consideration is that it, so it is memory consuming. Additionally, the authors consider a supervised pre-training stage for the convolutional network.

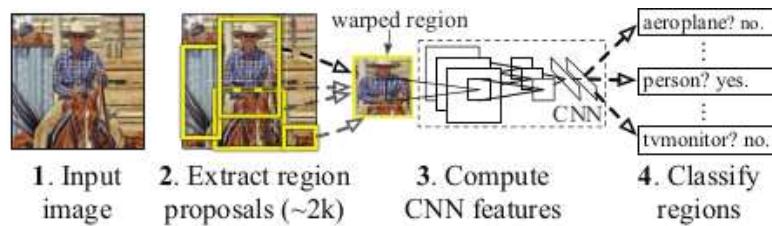


Figure 5. Basic R-CNN algorithm. [25]

R-CNN architecture appeared as an innovative method, combining the approaches of region proposal and convolutional networks to produce the bounding boxes. This result is a promising approach if disadvantages could be solved, thus developing in the study have been done generating improved versions of this architecture. The following subsections are devoted to discussing the most popular R-CNN based methodologies presented in the literature, that solved the mentioned issues. Moreover, we will discuss architectures similar in purpose but different in nature, such as the renowned You Only Look Once (YOLO) architecture.

2.3.2. Fast R-CNN

Similar to the standard R-CNN, this method inputs regions of image into the convolutional layers, but at the same time inputs the whole image. Besides, it combines different architectures to process the data such as ConvNet (a sequence of convolutional layers to pre-process the input image), Region of Interest (RoI) pooling layer for region proposal, and a classification layer. These architectures reside the same principal one, as shown in the following figure [25].

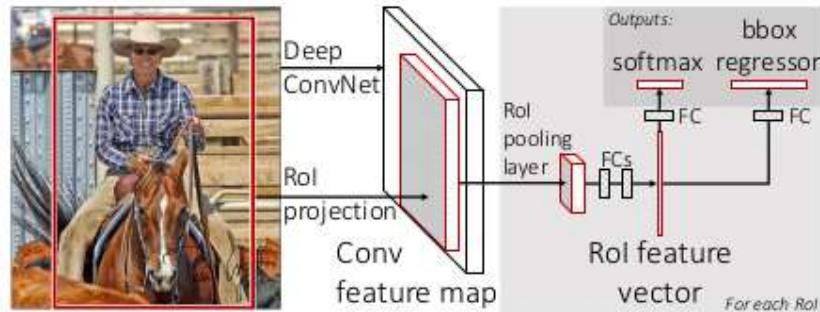


Figure 6. Fast R-CNN architecture [25].

This architecture dramatically improves the training time, over ten times the normal R-CNN, and achieves higher precision. Another contribution is that no storage is required for feature acquisition, reducing the memory needed. Although it drastically improves the time to about 2.5 seconds, it is still considerably slow for real-time applications; the selective search region proposal generation algorithm takes the most of the time. A newer architecture called Faster R-CNN attempted to overcome the problem mentioned above.

2.3.3. Faster R-CNN

With aims to real-time object detection, Faster R-CNN is a consolidated algorithm proved of capable, based on a Region Proposal Network (RPN). RPNs starts by taking an image as input. As output, they provide a set of different rectangles of object proposals. In this architecture, the selective search algorithm is different with respect to the previous R-CNN architectures. RPNs share full-image convolutional features with the detection network. It is a fully convolutional network (FCN) that simultaneously predicts object bounds at each position. This network is trained in an end-to-end base to learn the task of generating detection proposals. First, it generates a convolution feature map, which then can be passed into the RPN. Finally, it is passed into a classification layer responsible for determining if there is an object or not [25]. The mentioned process can be visualized in the next figure (obtained from [25]).

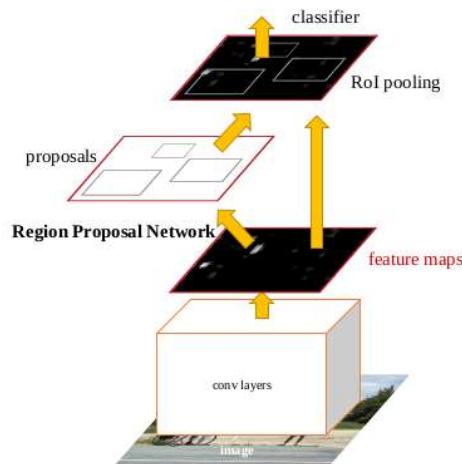


Figure 7. Faster R-CNN algorithm [25].

As the original paper's title proposes, this architecture's main objective was to get as close to a real-time object detection algorithm as possible. This architecture has achieved outstanding performance

and time response (between 0.05 and 0.2 seconds). The proposal's accuracy increased critically in various applications, being its most significant advantage than other analogue methods. It uses a reduced number of needed rectangle region proposals to 300, minimizing the computer processing need. One drawback compared to its most similar architecture (i.e. YOLO) is that all R-CNN architectures are two-stage algorithms, with YOLO being a single-stage one. Hence, YOLO is faster. YOLO is proposed as the faster approach in the literature, while Faster R-CNN is more accurate. The YOLO algorithm will be explained later in this same section.

2.3.4. Mask R-CNN

This method extends the Faster R-CNN concept by adding a branch for predicting an object segmentation mask in parallel with the bounding box prediction. In that case, the two units work for class prediction, both within an ROI. Segmentation is widely used for estimating poses, although it requires a much more refined spatial layout of an object. This method harnesses the velocity and principally the accuracy of faster R-CNN to get over the precision requirements and get over the known image segmentation approaches. [25].

A mask encodes an input object spatial layout, and thus, a mask is predicted for each ROI using an FCN. This mask maintains the spatial layout without collapsing into vector representations. Instance segmentation is challenging since it requires correct object detection and precise segmentation to classify each pixel into a fixed set of categories. Recalling, this method adds a segmentation mask on each ROI. The added branch for segmentation comprises a small FCN applied to each ROI, predicting a segmentation mask in a pixel-to-pixel manner. This scenario only adds a small computational extra requirement, non-considerable compared to the needs of the standard R-CNN methods, and other segmentation algorithms.

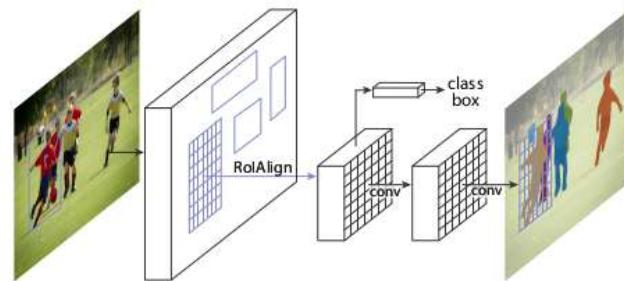


Figure 8. Mask RCNN architecture [25].

Furthermore, it is essential to remark that the authors proposed the RoI-Align layer, a quantization free layer instead of RoI-Pooling that has misalignments between the ROI and extracted features. ROI-Pooling does first the same step for all ROI abstraction methods, divides each coordinate by K , from the new relative coordinates the ROI is cropped. Finally, with quantization, the cut part is divided into bins that give an $n \times n$ ROI. RoI-Align presents a standard operation to extract small feature maps, properly align with the input to present the per-pixel spatial correspondence, by defining the ROI with 4 points from the feature's division coordinates map, which is then bilinear interpolated to result in the final ROI. This process results critical for the mask prediction, and has an enormous impact, improving the mask accuracy by 10-50%.

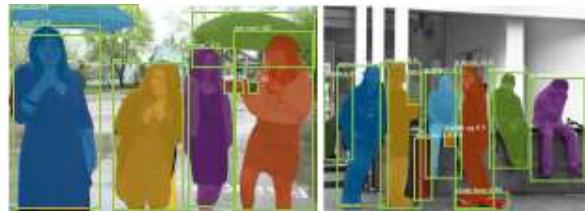


Figure 9. Segmentation results of Mask R-CNN [25].

The model outperformed the baseline of other state of the art instance segmentation models. Authors reported the inference time of architecture, capable of running at 195 ms per image on Nvidia Tesla M40 (plus 15ms CPU by resizing outputs). This method has been satisfactorily used in different applications. However, it presents a slightly higher accuracy than other mature architectures; there is the dilemma because there are faster architectures capable of real-time working such as YOLO (explained in the next section), which shows an adequate working velocity.

2.3.5. YOLO

YOLO deserves an analysis of its own, as it happens to be one of the most used architectures to perform recognition tasks. YOLO is an algorithm that uses a deep NN architecture output for making bounding boxes around all the objects that the algorithm recognizes in an image. The actual architecture used in YOLO is called DarkNet [26]. First, the image is divided into a grid, and in the last convolution layer of DarkNet, a sliding window analyzes each of the grid's dimensions at once. DarkNet uses a VGG-like [27] architecture with leaky ReLU activations in each layer; but in contrast to other architectures, has no fully-connected layers at the end. Instead, it uses a $5 + C$ vector. That is, concatenated to a C channel vectors, being 5 the key components for detecting all objects in an image, i.e. object presence, objects localization in x and y coordinates, and objects dimensions, and C corresponding to the number of classes that DarkNet is trained on (20 in the original paper [27]). This vector is then convolved into a final net output with the width and height of the grid dimensions chosen for a specific image. The last dimension uses an arrangement, where B is the number of boxes per grid that the net is capable of handling. Once this last layer is done, then the YOLO algorithm works its task by sliding a window through each pixel of the final convolution layer, which then traduces into analyzing each of the cell grids of the original image at once, hence the name of the architecture. Once a simultaneous sliding window in each grid cell is produced, YOLO searches for an object that DarkNet can recognize, derived from the training data. In the original paper, each grid cell has two boxes, meaning that $B = 2$, giving a channel dimension of 30. These boxes exist because it could happen that two objects share a single grid cell, and so because there are two boxes, then two bounding boxes can be drawn and recognize both objects. B can increase so more objects are recognized per grid cell. In further YOLO applications B is commonly set to 5. This makes YOLO a very fast and accurate algorithm, being very suiting for real-time applications.

Despite this, YOLO doesn't cope very well in very crowded scenarios. Given the grid approach, a minimum-maximum number of elements can be found in each grade depending on the setting. To address this issue, the last layer would be needed to be retrained with a finer grid and more boxes per grid cell, thus complicating the architecture, as not only YOLO would need to detect more classes, but also needs to be retrained to see more objects of the same classes it has already been trained for. It is important to state that YOLO has been experiencing several updates which lead to improved versions. For instance, YOLO v4 is the best one, exceeding the 40 FPS in and NVIDIA GTX 1080 Ti GPU architecture, having a 65.7% mAP. It uses only three boxes for the B variable. Its training uses weighted-residual connections, cross mini-batch normalizations, and other techniques that make YOLO v4 very fast and accurate compared to different object detection algorithms.

2.4. RNN

A Recurrent Neural Network (RNN) is another type of architecture designed to model time series. A multilayer perceptron has the distinction of allowing connections among the hidden layers to describe a time delay. This architecture is known as the short-term memory of the NN. Because of the kind of relationships mentioned before, the model can retain information about past states, allowing to understand temporal correlations between events distant in time to each other [28].

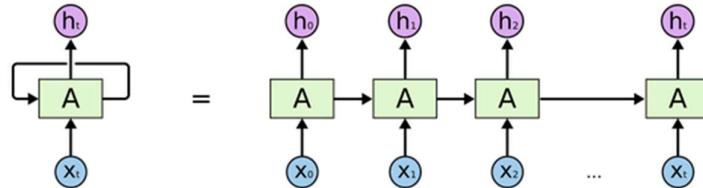


Figure 10. RNN Architecture [29].

Its principle, RNN presents a robust and straightforward solution to predicting possible future scenarios based on previous events allowing information to persist. Its initial architecture presents two main problems at the update stage of backpropagation; The *vanishing gradient problem* and the *exploding problem*. Both stand for improper training or wrong updating. On the one hand, the vanishing gradient describes that the neurons at the very back take too much time to update or even stop being updated. Nonetheless, the exploding gradient problem represents the closest to present neurons, which tend to have the highest weights, which gives less importance to farther weights. Both issues have been solved in more advanced architectures derived from the classical RNNs. RNNs are brought to this work because they have been used as an extra tool for object detection, taking advantage of its memory capacity; this can be seen in the next section.

2.5. CRNN

The Convolutional Recurrent Neural Network (CRNN) is the combination of the two previously mentioned kind of NNs; the CNN (Section 2.2) and the RNN (Section 2.4). The CRNN architecture has been widely used in several applications such as sound recognition and classification [30], [31], where they have found particular success.

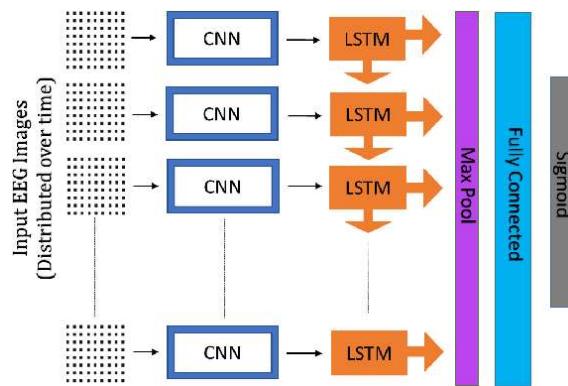


Figure 11. CRNN Architecture, in which a CNN and a type of RNN are present among other layers [32].

CRNN based methods were developed to generate more profound and accurate image analysis. Given that convolutional layers have mainly been used to obtain image representations and features, these methods were enough to get acceptable results in many cases. However, as the need for a better

quality of analysis has increased (due to the high demand for accurate classifications), recurrent layers have become a promising solution in combinations with the ones mentioned above. This setting offers the possibility to remember and connect independent features.

Zuo et al. [33] proposed a CRNN model consisting of a CNNs as a first step to extract middle-level image features and general image analysis. In series to the convolutional layers, the data must be adapted to suit the recurrent layers; this can be done by converting the regions previously obtained into a matrix of the total amount of areas. After this step, the data can be analyzed in a parallel recurrent layer of 1-dimension for each one. From this output, dependencies and correlations are generated from the recurrent layers, taking advantage of the memory capacity. Finally, the parallel layers must converge in a final layer that will merge all the gathered information.

This method's improvement was presented by Hu et al. [34], doing the same procedure of starting with a convolutional layer to extract image features and then feeding the data into the recurrent layers. This method proposes the implementation of the CRNN module in between any of the convolutional layers. What differentiates this module from a classic RNN is that the connection between layer and layer is convolutional and not fully connected. This kind of approach demonstrated better performance, reliability, and accuracy, although they consume more computer capacity for each of the modules used.

2.6. LSTM

Long Short-Term Memory (LSTM) is a recurrent-type cell for regressive NN. LSTMs are explicitly designed to avoid long-term dependency. This architecture consists of three gates, which are mainly intended to solve gradient disappearance and explosion in the long sequence training process. Put merely, LSTM can perform better on longer sequences than normal RNNs.

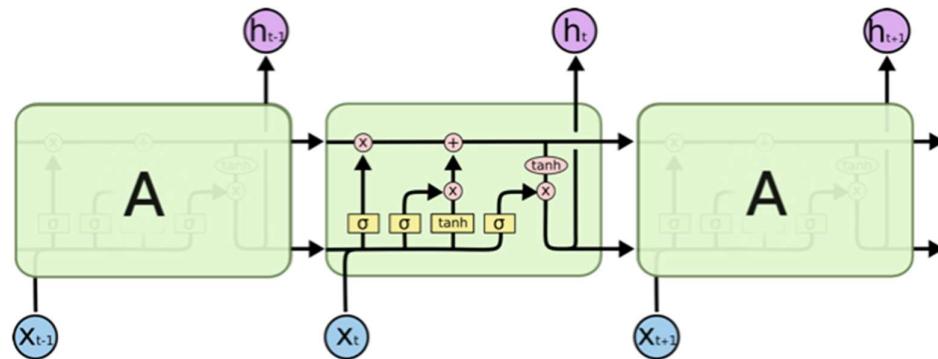


Figure 12. LSTM architecture, with three gates composed of four NN layers [35].

The previous image shows in more detail the structure of an LSTM, which, contrary to the standard RNN, the repeating modules has four neural network layers instead of a single \tanh layer. The principal idea behind the LSTMs is to have the ability to add or remove information to the cell state, which is the horizontal line running through the top of the diagram, where the data could flow along with it unchanged. This ability is regulated by structures called gates composed by a sigmoid neural net layer and a pointwise multiplication operation. Depending on the digital output state, the information would be let through. LSTMs have three of these gates. The first one decides which inputs are going to be ignored by the cell state. The second one determines which values will be updated. With the help of a \tanh layer, which creates a vector of new candidate values, an update to the state will be produced to decide what new information will be stored in the cell state. Finally, the third one decides what parts of the cell state will be the output [35].

2.7 GAN

A Generative Adversarial Network (GAN) is a type of NN that consists of two architectures: a discriminator D and a generator G . In principle, the idea is that based on a minor amount of training data Z , D is capable of generating a set of artificial data X_{fake} . Afterwards, D is trained on X_{fake} and real data X_{real} to recognize the difference by estimating losses from prediction and ground truth. Notice that in contrast, G estimates the log-likelihood of the distributions of the data. The early version of this architecture, often called Vanilla GAN [36], uses the *Kullback-Leibler* divergence as the distribution similarity to produce the false images. Some other architectures presented in the literature are C-GAN [37], S-GAN [38], AC-GAN [39] and F-SGAN [40], with an ever-growing market in the research community⁷

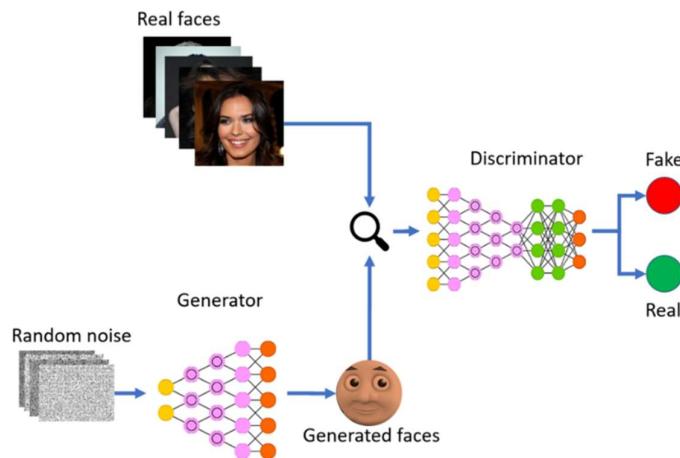


Figure 13. GAN architecture [41]

These networks are typically used to generate more data samples so that the data used to train classification systems can be improved. Most notably, this has been applied for industrial applications such as face detection [42], fish classification [43] and symbol recognition in engineering drawings [44].

2.8 Distributed Systems

According to Coulouris et al. [45], the concept of a DS refers to different agents connected working for a common system. Each node communicates through messages to achieve a common goal. A DS in the cloud has the following characteristics:

- *Concurrency*: Each node can perform different tasks or parts of the same task-sharing resources, allowing parallel computing.
- *Quality of service*: this concept refers to the idea that a specific type of message's sensibility depends on the type of information transmitted. For example, if you download a file, it may not be vital if it delays for some seconds. However, in a phone call, or a video call, any delay may result in a loss of data that troubles communication.

⁷ Hindupuravinash, 'The GAN zoo', *Github*, [hindupuravinash/the-gan-zoo: A list of all named GANs! \(github.com\)](https://github.com/hindupuravinash/the-gan-zoo)

- *Scalability*: Any network must be built having in mind that it may expand.
- *Security*: With the constant growth of devices connected to the cloud, and the high amount of information transmitted, safety must be implemented on every network designed.
- *Failure handling*: In a DS, there must be a plan for possible failures. If a computer fails, the others will keep working.
- *No global clock*: This happens because of the communication through messages. Each node may have its frequency of operation as each one has its own hardware. This situation may difficult the transmission because there is a limit to computer synchronization accuracy in a network. Therefore, a bottleneck in communication, where the best bandwidth that may be reached is the lowest in the network, may be caused. However, this allows each node to operate independently with its requirements of hardware.

However, it's important to note that each node is a different device because, in edge computing, they will process the information without sending it to the servers. The nodes that are closest to the sensors can be the ones that process the information because they will have a processing unit, memory, and the hardware needed resulting in less latency, less energy consumption, more privacy, and a faster process overall. When deep NN grow in complexity, usually the training time increases as well. CNNs, as well as other architectures mentioned above, might face this challenge, especially when several convolutional layers are required. With the current average computer systems, a CNN might end up needing a couple of hours to train, not to mention the time it takes to process the input data, along with any other specific operation.

Due to this hindrance, several attempts in the literature to develop DS applied on NNs, now on referred to as Distributed Deep Neural Network (DDNN), have been presented. DDNNs implemented through DS offer the same advantages with regards to the ones mentioned above for DSs. In [45] and [46], the authors modified the raw data before sending it to the NN in the cloud. This prototype is an example of edge computing applied in DDNN to avoid bandwidth limitations in network communications. Also, this helps with training time, because the information will be filtered. As a result, the system's privacy is strengthened because all information packets will only be available locally on the nodes that gather data. It is expected that, similar to image preprocessing, NNs in the cloud need to modify data before feeding it to the NN. One might even say it is necessary because of the limitation of bandwidth. Usually, gathered data needs to be preprocessed before uploading it to the cloud as input to the NN, as shown in [45] and [46].

Usman et al. [46] implemented a cloud-based video analytics system. They used several layers in the system to process the information. They used parallel model training, which trains several partial models and combines them in a central one. The information received was filtered by separating consecutive frames that lacked a significant change in the object. These characteristics reduced resource usage and training time by filtering video frames that lacked an essential change in the image observed.

Teerapittayanon et al. [47] proposed a DDNN, which works over a cloud, edge, and end devices. This system was used to analyze different sensor data, minimizing resource usage for devices by processing a portion of the information recovered in the end devices, and an edge layer. This separation also allowed system fault tolerance, and enhanced data privacy, because each end device didn't offload raw sensor data to the cloud directly.

Although this approach may come with several advantages, it also brings a handful of problems to the table. Unlike centralized systems, information transferred faces the problem of bandwidth. Data

cannot travel at the same speed. In Internet of Things (IoT) systems, the filter can be local, right before traveling to the cloud. This way, useless or repetitive information can be disposed to reduce time and space for training, allowing the part of the DS in which the NN is present to focus all its resources to train on sufficient data. In [46], a CNN was trained to classify objects. Although videos were set to 25 frames per second, the system only used five frames each second.

2.9. Metrics for evaluation

Several metrics exist to evaluate the performance of object detection and segmentation tasks. These are useful to assess the behavior of a new model and to compare it to existing ones. Furthermore, there are competitions, such as the Common Objects in Context Object Detection Challenge (COCO-ODC) or the PASCAL Visual Object Classes (VOC) Challenge [48] that use a combination of metrics to serve as benchmarks of quality for object detection models. Before explaining said metrics, some concepts need to be defined. Firstly, the Intersection over Union (IoU) or Jaccard Index, is defined as the intersection between the predicted bounding box and the ground truth bounding box over the union of those two, or simply put, the similarity between the prediction and the ground truth. This score ranges from 0 to 1; the higher the score, the closer the prediction is to the ground truth. This metric can generate four outcomes: 1) True Positive (TP) the model detects an object where, indeed, an object exists, 2) False Positive (FP) the model detects an object where there is no object, 3) False Negative (FN) the model detects no object where there is an object present and 4) True Negative (TN) the model detects no object where there is no object. The next important concepts are precision and recall. Precision refers to the probability of the prediction matching the ground truth. It can be defined as $\frac{TP}{TP+FP}$ or $\frac{P}{ALL\ DETECTED\ BOXES}$ [49], [50]. In counterpart, recall refers to the rate of TP or the probability of a TP to be detected. The recall can be defined as $\frac{TP}{TP+FN}$ or $\frac{P}{ALL\ GROUND\ TRUTH\ BOXES}$. Adding on, the precision-recall curve is a metric that summarizes the two previous concepts and is especially useful to assess the performance of models with imbalanced datasets [30]. In this curve, precision is plotted on the $-$ axis and recall on the x -axis. Ideally, a model should deliver both high precision and recall, although this is not always the case. When comparing models, generally, the one with the highest curve indicates the best performance. Another way to gauge these two metrics is the F-measure or F1-score, defined as $F1 = 2 \times \frac{precision \times recall}{precision + r}$.

Furthermore, it is also possible to calculate average precision (AP), which, as the name indicates, is the average precision among the precision-recall curve. Mean average precision (mAP), on the other hand, involves the AP among several classes, whereas the AP only covers one. The mAP is defined as $mAP = \frac{\sum_{i=1}^K AP_i}{K}$. Regarding the competitions mentioned above, the COCO-ODC uses 12 metrics, where mAP is the main one. On the other hand, PASCAL VOC uses two metrics: the precision-recall curve and AP. It is also worth mentioning that different challenges use different datasets with different scenarios in mind. For example, the KITTI dataset consists of hundreds of hours of driving systems gathered from an on-board camera on a car, which is especially useful when working with AVs.

On the other hand, the COCO-ODC dataset contains many frequently found objects (80 different classes) such as animals, tools, people, etc. The PASCAL VOC data set is very similar, and both are commonly used when training object detection architectures. However, there are some key differences worth noting. For instance, annotations in COCO are facilitated by employing a JSON file, while PASCAL uses the XML format. Furthermore, the bounding box format is different: COCO uses x_{min} top-left, y_{min} top-left, width, height, while PASCAL uses x_{min} top-left, y_{min} top-left, x_{max} bottom-right, y_{max} bottom-right [48]. Although these differences may sound minimal, some researchers decided to focus on specific challenges based on the formats that they are more familiar with. Thus, there are no common grounds to compare every single algorithm ever made.

2.10. Frenét Motion Planning

The Frenét frame method consists in transforming the Cartesian coordinates x and y to a more dynamic reference frame n and t , which are normal and tangential to the road vectors, respectively [51]. If a motion planning algorithm in a highway uses the traditional cartesian coordinate plane, it would have to parametrize x and y coordinates positions concerning time and solve for polynomials with significant coefficients. This is a challenging and computationally expensive task. Thus, the Frenét coordinate system uses vectors projected on the road's curves, treating the road as if it was all forward plane with no curves. Hence, by changing these vectors magnitude, lane changing and braking and acceleration happen to be an isolated problem from the road following task. It helps to have two different easier problems: lane keeping and lane changing; instead of one more difficult situation. We favor the use of Frenét coordinate systems since these allow for lane-keeping using minimal computing effort. They also predict and calculate lane deviation in case of dodging for static or dynamic obstacles, to then return to the lane's boundaries. The Frenét algorithm does not execute the prediction task, but its output can be easily used in many prediction algorithms. The Frenét coordinate system can also be taken for peripheral vehicles other than the ego. Still, its original purpose remains for the car to perform inertial odometry tasks. Much more can be exploited, but to get it to work, many sensor fusion tasks need to be performed to state the vehicle's center of mass, since it will be the origin for the Frenét coordinate system. For example, in Figure 14 a route in which x and y coordinates are the curve's points, the vehicle doesn't see those coordinates. Instead, it transforms its coordinates to vectors T and N , so the car perceives itself as traveling in a straight line, only changing N vector to keep the lane and vector T changing its speed. To update its N and T vector values, the following linear algebra equation is used: $\begin{bmatrix} T' \\ N' \end{bmatrix} = \begin{bmatrix} 0 & \delta \\ -\delta & 0 \end{bmatrix} \begin{bmatrix} T \\ N \end{bmatrix}$

Since each car's Frenét route needs to be predicted, the T and N vectors are given by combining many algorithms, as discussed later in Section 3. The T and N derivatives are calculated with respect to the arc length, and thus the only missing value, and the one that needs to be predicted, is δ . Once δ is predicted, the Frenét Frame can be calculated by solving for the x and y curve which is tangent to T and normal to N . The updated frame for each point of δ will be 35 Hz as it is the vision algorithm frame rate and T and N will be captured using computer vision algorithms.

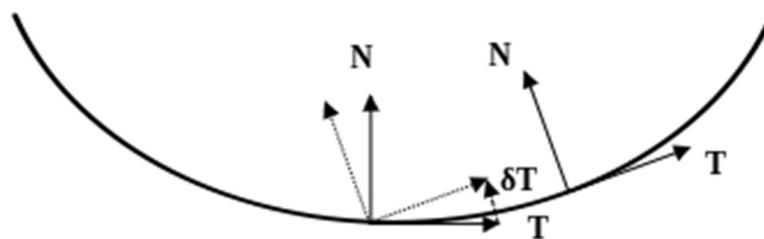


Figure 14. Frenét frame visualisation [52].

It is essential to state that Frenét frame is mainly used for motion and path planning, and so the equation given above is not the only one.

3. Applications for AVs and Assisted Driving

This section is divided into four main categories. The first one presents image pre-processing and instance segmentation, in which the main aim is to locate all *legal* boundaries and other objects of reference. It is inferred that although one vehicle is physically capable of driving on a sidewalk, it is not correct for it to do so, and therefore the ADS needs to recognize the road lanes. Nonetheless, not

all road lanes are used for driving on the required direction (i.e. single direction only, cycling/pedestrian lanes). Thus, this step plays a crucial role in the detection of which road lanes shall the vehicle use. The second one focuses on object detection within the semantic segmenting boundaries of the route. For instance, algorithms such as YOLO are used to recognize each object, and this data is fused with LiDAR point clouds using an MV3D-like architecture [12]. Processes such as these help the vehicle draw prism points on each detected object and know where its center of mass relies upon, and the volume being handled within the lane space. Human and sentiment detection techniques (for driver, passengers and pedestrians) are discussed in the third section. Finally, in the fourth category, we discuss cloud computing and parallelization to speed up and improve the pipeline. For instance, these techniques can train all object detection models, communication of different vehicles, etc.

3.1. Image Pre-Processing and Instance Segmentation

Chen et al. [17] proposed a CNN model for object localization using camera and LiDAR data for 3D multi-view object detection, or MV3D. In this method, three CNNs are deployed in parallel. Two of them are fed with the point cloud LiDAR data (one obtained from a BEV and the other from the front view). The third CNN uses raw RGB images from a camera located at the front, on the sides, and the car's back for 360 view degree. Once each architecture finishes, a deconvolution process occurs to identify where the tracks are located and their shapes. Like an RCNN architecture, a decoder-like function outputs decoding fed to the architecture performing the localization task and shared between the other parallel architectures. This process generates a 3D object proposal of the classes, combining a region-wise feature, obtained via ROI pooling caused by each architecture. The output is a 3D bounding box around the recognized objects. It is important to note that this method can realize the item's class and position and its volume. This method, performs with a 14.9% higher accuracy than state of the art methods, being comparable in runtime. As a result, this method is computationally demanding, and deducting the 3D bounding box will require a specific and robust hardware architecture for the algorithms to be deployed correctly and in real-time scenarios.

Zhou et al. [52] proposed VoxelNet, an end-to-end trainable deep architecture for point cloud-based 3D detection. In contrast to the previous two methods, this is a LiDAR-only based detection, which relies on the object size, occlusion state, and other superficial information. To improve the localization and detection, the authors proposed the Multimodal Voxel for 3D Object Detection Network (MVX-Net) framework. In this method, fused augmented LiDAR point clouds with semantic image features are used at early stages to improve 3D object detection. Moreover, two fusion techniques were developed to extend the VoxelNet. The first one is PointFusion, which combines features at early stages to add 3D points to an image feature and capture dense context, extracting features from a 2D detector (i.e. a camera) and then jointly inputting them into the VoxelNet model. The second technique is called VoxelFusion and employs non-empty 3D voxels to be projected on the image plane. Features are extracted and fed into the voxel feature encoder used by a 3D RPN to deduce the bounding boxes. Authors state that the VoxelFusion technique delivers a slightly inferior performance compared to the PointFusion one; however, it has more efficiency in memory consumption. Authors evaluated the model on the KITTI 3D object detection benchmark and concluded that their approach demonstrated better results than the single modality VoxelNet [53].

Avoiding person-to-person contact is the best way to control and prevent the virus' transmission amidst the COVID-19 epidemic that affected the world in 2020. Still, people require food and other basic goods to address their essential needs. Liu et al. [53] developed Hercules, an autonomous logistic vehicle used for contactless delivery to address this issue. It has an autonomous navigation capability and a stable hardware-software system to manage the operations. Its maximum payload and capacity are 1000 kg and 3m3, sufficient for most of the people delivery's requirements, and uses

3D object detection with LiDAR point clouds to recognize and classify objects. Multiple calibrated LiDARs provide data which are fused as input into the 3D object detector of the VoxelNet. Authors consider that the real-time performance deserved much more attention than accuracy for this particular application. To improve efficiency, they replaced dense convolutional layers with spatially sparse convolutions to obtain inference time boosting. Another critical part for Hercules' perception task was to build the 3D map of the environment with 3D point clouds from the LiDARs and the readings from the inertial measurement unit. Objects in motion should be treated differently than static entities. Thus, this task is very crucial for motion planning in AVs. Hercules' solution is called Ego-motion Compensated Moving-object Detection, which detects objects based on the consecutive point cloud frames. The autonomous navigation technology was reported to work well on state of the art scenarios [54].

The Intelligent Autonomous Robotic Automobile (IARA) was presented in Moraes et al. [55] and Cardoso et al. [56]. It is a vehicle that uses a suite of sensors (including a LiDAR and odometry sensors) to perform various tasks. The two papers present a NN aided approach to state of the art navigation, mapping, and localization methods. The first publication presents a novel, image-based real-time path planner. Using a CNN, the algorithm can deduce a path from images the on-board camera sends. This path (a cubic spline) is then transformed to the car's coordinates, and thus the route can be followed. The second publication presents a new approach for real-time inference of occupancy maps, using deep learning. The network proposed in this paper, NeuralMapper, takes the LiDAR data as input and creates an occupancy grid map. It further uses a NN to infer certain parts of the route where LiDAR data is not very accurate, creating a completer and more efficient map.

If a car detects an obstacle, ~~they~~it brakes and corrects ~~their~~its course, but ~~they~~it does so when the interruption has already occurred. Analyzing human gestures can help avoid sudden collisions; however, deep learning models do not have enough data to train them. For this reason, Cruise's AVs [18] use motion capture to understand human gestures technology, which is a technique that video game developers use to create and animate characters. With this technology, the necessary characteristics would be extracted to train the deep learning models and be one step closer to predicting the sudden movements of pedestrians or drivers.

There are two types of mo-cap systems optical and non-optical. The visual uses cameras distributed over a sizeable grid-like structure that surrounds a stage; the video streams from these cameras can triangulate the 3D positions of visible markers on a full-body suit worn by an actor. In a non-optical way, this system [57] uses a sensor-based version of motion capture instead, which relies on microelectromechanical systems (MEMS) [58], [59], which are portable, wireless, and do not require dedicated studio space. That gives a lot of flexibility and allows to take it out of the studio and collect real data of world locations.

Barea et al. presented an integrated state of art CNNs for multisensory 3D vehicle detection in a real autonomous driving environment [60]. As the title expresses, the proposal of their paper aims to present an outstanding architecture based on the combination of state of the art methods for object detection such as YOLO and Mask R-CNN for 3D segmentation besides of a LiDAR point cloud. Figure 15 shows the results achieved in this paper, integrating Mask R-CNN, YOLOv3 and cloud point LiDAR. It can be seen that an accurate bird eye view, which corresponded with the object detection in the superior image, is created. Moreover, an accurate mapping of the road's current state, with the vital information of the 3D sizing of the moving and static objects in the ego car's peripheric view, is obtained with complete details for the navigation. The unique architectures applied in this proposal achieved high performance in the KITTY test set.



Figure 15. LiDAR's bird eye view plus boxing object detection [60].

This work proves that individually each architecture cannot overcome themselves. Still, in the correct combination, new opportunities might be found. For the reader, in machine learning and similar technologies, the novel results are achieved based on experimenting and not in a fixed formula. Thus, the following question arises: What would be the best set combination for acquiring the correct data for a completely autonomous car?

As mentioned previously, autonomous driving is a challenging problem that has been addressed with various technologies, methods and algorithms. Still, it seems that there's always space for more efforts, due to the rigorous task of taking decisions in an always-changing and a mostly unpredictable environment, especially on busy streets and highways. The following paper [61] discusses how image segmentation allows the class to classify the road's drivable and non-drivable regions. To do so, a Mask R-CNN architecture was trained to differentiate these possibilities. The architecture was trained and tested in the Berkeley Deep Drive (BDD100k) dataset, with 100k on road pictures, achieving a final mAP of 0.79. To identify the drivable areas, the model determines the ego car's current line, making the difference with the alternatives. The implementation problem was broken in various subproblems; road object detection to declare a line as non-drivable, segmentation breaking a 2D Image into depth-based layers, and lane marker detector, to separate each of the lanes.



Figure 16. Line segmentation based on Mask R-CNN [60].

This paper also explores new possibilities demonstrating that the architectures mentioned in the previous part of this paper can be used in different ways to achieve similar objectives, which provides more information to the ego car, to be able to make better decisions by evaluating the state of the environment. On the other hand, this paper doesn't mention the time it takes to the algorithm to drop a given state's results.

3.2. Driver Assistance and Predicting Driver Patterns

One characteristic that differentiates a good driver from a bad one is the capability of choosing the best possible action, in the proper time, and execute it correctly, at all times, even in the rarest and most dangerous circumstances. In those cases, being a skilled driver can save lives. In the context of AVs, this is a real threat because, as we have mentioned in various occasions, one of the biggest challenges for self-driving cars, is the unpredictability of the always-changing road conditions and human behaviors. At its beginnings, AVs were based on precarious algorithms, based on state conditionals, with the first driver-assisted systems, acceleration and deceleration, and conditional automation, where human intervenes are expected when needed. These old approaches were not able to counter the odds of real-time decisions⁸.

Nowadays, state of the art aims to evolve to fully automated systems, where non-human intervention is needed. But one barrier that challenges all done developments is how a human perceives the driving and, consequently, decides and how a computer does. Many background contexts conditions are known for humans that can be subjective and for computers becomes complicated to understand; for example, the local cultural manners of driving often define what to expect from other cars, which differ in geographic location. Another condition is the passengers' feel. This concept means that it might be the same for a computer to avoid an object a second before the collision; therefore, this move can be made smoothly. Conversely, this decision for the passengers may seem uncomfortable, causing other issues, such as a human hitting the window. Therefore, the autonomous system should consider the passengers' inner motions, thus being dynamic, and based on human behavior to improve it, without realizing actions that might confront what a human expects from another car.

The mentioned considerations are causes and circumstances that must be considered for a full automotive driving system. Thus, the decision-making part of an autonomous system is a critical part

⁸ WIRED Brand Lab, *A brief history of Autonomous Vehicle Technology*, WIRED, 2019, [A Brief History of Autonomous Vehicle Technology | WIRED](#)

of the process. Driving patterns are ways in which a human driver behaves. These patterns are developed through experience and practice. Inside these driving patterns, humans can predict how the environment can change and based on these predictions, and a decision is made. There are a wide variety of methods and solutions that make use of different NN architectures [62]. Some of them are:

3.2.1. Behavioral Cloning (BC)

BC consists of a CNN that uses the real driving data obtained through the vehicle's camera and computer when driving during different situations. Then, the network learns the response of the driver to obstacles in other locations. According to [63], a whole model can be trained using only raw data and expect to have a steering output. This approach is called end-to-end learning, but electromobility and autonomous driving named BCBehavioral Cloning are much more appealing. End-to-end learning is called that way because there are no subprocesses when expecting and output from the NN. These subprocesses include SLAM, Path planning, object recognition, etc. None of these techniques is used for motion planning of the vehicle; the only information used is data taken from feature extraction in images taken when the car was being driven. These images have the steering angle and gas pedal position as a label for the dataset. Thus, when seeing a similar scene to the ones in its training data set, then the gas and steering values are close to the ones it was trained on.

It could be very promising that there are no subprocesses required for achieving autonomous driving. Using raw images, then the only two values needed for driving are the ones generated. But the reality is much more discouraging. It would require enormous amounts of data, which is not at hand. Because of the model generalization, millions of different driving scenarios need to be recorded but have not yet been taken. BC works only during certain situations, but it works encouragingly. BC is a method that deserves to be discussed in the future once more data is available. For the time being, end-to-end learning is used and applied in natural language processing because of its capability of finding sense in sentences, compared to traditional methods. Data in language is large enough for end-to-end learning be applicable, not the case for autonomous driving in which research is being developed but years away from practical applications.

3.2.2. Reinforcement Learning (RL)

One of the most famous machine learning paradigms is RL; this type of algorithm has been used in different ways to process the information given by the sensors to output the best actions. This broad field of study aims to replicate with algorithms how living beings learn through the process of gaining experience by exploring its environment and realizing which activities are suitable and which are not, based on rewards, in relation with a specific objective or task inherent to the environment. RL has been widely used as a control algorithm to determine the best policy, or in a simple word a strategy, that will determine the best possible action for any presented state or circumstances that the environment presents [64]. Figure 17 shows the basic flow of this type of algorithms.

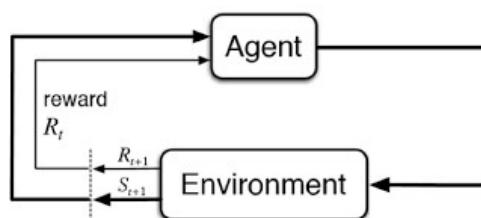


Figure 17. Basic RL process, agent-environment relation [64].

The necessary process for an RL to work is as follows. An agent that is the learning object interacts with the environment, which can be virtual or real, the agent performs actions (a) given an environment state s , that will affect the environment, in exchange, the environment will feedback the agent with a new state s' and a reward r . If the action were right, then the reward would be positive, contrarily if the move were wrong, the reward would be negative; in other words, a punishment. The purpose of the agent is to maximize rewards. The process aims to create a system that behaves in the best possible way to achieve a goal.

With a good overview of RL and without deepening the algorithm's details because they are out of the scope of this work, we can now discuss how this paradigm has been applied to AVs. This type of algorithms is used in the second part of the process for the vehicle's decisions process. First, it has to acquire data of the environment, in this case, the road and the objects such as other cars interacting within it, via the already mentioned and studied approaches such as LiDAR and object detection algorithms via image processing. Based on these data, advanced RL versions based on NNs, such as deep RL, deep Q-learning, and deep deterministic policy gradient (amongst others), have been used for AVs' decision-making processes. These deep versions of RL present the advantage of a more robust, more powerful capability of learning behaviors, understanding patterns and predicting situations in comparison with their non-deep versions.

An example is control of speed and lane change decisions [65]. As the authors mention, a known problem of decision taking on AVs is that the methods are designed for a specific task. An example would be the Intelligent Driver Model (IDM), which is applied to decide only when to change lines. Therefore, there's a need for a more global algorithm to take care of the multiple decision needs for AVs. This work proposes that the car uses a 27-input vector providing information on the road and eight possible surrounding cars to make lane change decisions and speed control. However, in the same DQL algorithm, they trained two agents to address each of the tasks [66].

One perceived disadvantage of RL algorithms in autonomous cars is that they cannot predict the intentions of other road users, but just evaluate the current state of all the objects. These kinds of learning don't have long-term future recognition of the environment. This problem is intrinsic in the math that the algorithms are based on, because Markov decision process, take decisions based on last observations state.

3.2.3. LSTM based models

Currently, there are many approaches to road prediction using LSTM architectures for this type of forecasts. Bai et al. [67], suggested spatiotemporal LSTM architecture for motion planning in AVs. It uses convolutional layers for feature extraction in an image, and it runs through an LSTM to find sequences in those images, something similar to a CRNN architecture. Once this data is acquired, a CNN architecture is applied to extract spatiotemporal information from multiple frames. Lastly, a fully connected layer is used to extract the trajectory taken by peripheral vehicles. This trained model can be of service for detecting multiple object trajectories. It is considered an end-to-end algorithm since all this raw data enters this architecture, and it outputs a solution. Also, it is quite demanding in terms of training and testing data.

Atlch   et al. [68] propose using an LSTM architecture for lane changing detection on highways, using data such as velocity and lateral displacement from vehicles. It was trained using NGSIM US-101 dataset and trained two different LSTM models, one for lateral displacement and another for linear velocity. It performed very rapidly with 70 cm of error in displacement and 3 m/s error for speed, with 10 seconds ahead of raw predictions. This is a less direct approach since it requires the labeling of data. It can be assumed that they are using a Fren  t like topology since they are only considering linear velocity and lateral displacement in their model.

3.3. Human Sentiment Analysis

One of the most critical objectives for intelligent transportation applications is road safety, and unfortunately, one part of the traffic accidents is human error. A driver's behavior could affect its way of driving. For that reason, a monitoring system inside an intelligent vehicle would provide relevant information and indicate if the driver's action is allowed. Otherwise, the car would be able to correct it. This monitoring system used cameras inside the cabin of a vehicle to detect face sentiment. Then, CNNs are applied for the action recognition, to be more specific R-CNNs are the ones than select the most informative regions from the drivers [69].

These NNs are trained and tested with an extensive data set to detect any change in gaze and facial emotions, thus discerning regular and aggressive driving [70] or distracted driving actions [71]. In some other cases, there want to be able to predict drivers' behavior; to make this possible CRNNs are used. Kim et al. [72] proposed a line-segment feature analysis-convolutional recurrent neural network (LFA-CRNN) to analyze drivers' facial expressions indicating pain to prevent automobile accidents and respond to emergency health-risk situations. The LFA-CRNN showed an accuracy of 97.4% in contrast with 98.21% and 97.4% of CRNN and AlexNet.

Other systems also analyze people's tone voices inside the cabin, such as *Affectiva Automotive* [73], which understands what is happening inside a vehicle. This technology works as both a driver monitoring tool, ensuring that safety drivers keep their eyes on the road even as the self-driving software drives the car, while an emotional tracker ensures robot taxi passengers feel safe during the trip. The system monitors levels of driver fatigue or drowsiness, driver distraction, understand driver mood and reactions and could enhance fleet management. To solve these tasks, the software uses some deep learning architectures that have been mentioned in the previous sections. For instance, for the face detection, tracking and classification are used CNNs, in specific RPNs; and for the voice-activity detection and classification RNNs, in specific LSTM.

3.4. Using the Cloud with Autonomous Driving Systems

To solve the limitations posed by cloud systems, there are different solutions in the literature. In this section, we describe some of the solutions proposed. First, we need to explain how the cloud helps autonomous driving systems. According to Liu et al. [74], the cloud has to provide distributed computing, distributed storage, and heterogeneous computing. It is expected that these systems are tailored together. However, this creates problems with resource sharing, because the information has to be copied between each application, reducing available space and speed. In this paper, the authors propose a unified infrastructure that complies with the services mentioned. The system developed was reliable, had low latency, and a high-throughput autonomous driving cloud. These practices are useful if we want to build a system in which several AVs share their information gathered.

Kumar et al. [75] developed a cloud-assisted system that involved several vehicles communicating with each other and with several sensors located in the environment. To avoid bandwidth limitations, they used a request-based system, which allowed the car to have information about different locations at various resolutions. They used the Octree representation, which is a 3D-point cloud which can be divided recursively into eight. For more information, the reader is referred to [75].

4. Future Trends

Research on AVs has shown great promise, and there have been so many breakthroughs in recent years. However, there are still many challenges that the community needs to overcome to provide the safest and robust vehicles. All the sensors involved in ADS produce a lot of data from their surroundings, and they are required to be processed in real-time so that the AV can make a correct

decision; this is very critical because a small delay could make a considerable change. In this section, we will present some future trends in ADS, which can improve performance and reduce response times.

4.1. Cloud computing

Nowadays, cloud computing is becoming a trend for many aspects of our everyday life. More and more businesses are migrating several of their systems to the cloud because of its availability, security, scalability, among others. When discussing businesses' status, payments, or stock inventory, the amount of time required to get a result may be enough if a human perceives it as instantaneous, or even tolerated if a small delay occurs. However, in application domains such as autonomous driving, real-time results are a must if accidents or life-threatening scenario have to be avoided. Thus, the need to explore concepts such as edge, fog within the cloud computing domain. Edge and fog computing consist of moving the location where the information is processed closer to where it is generated. According to [76], any device with an internet connection which can store and process data can be used as a fog node. This way, the cloud server might rely on these nodes for most of the analysis reducing latency and increasing privacy. Authors considered three characteristics for an application that needs fog computing:

1. Data is gathered on an edge. In this case, an AV, or a group AVs can be considered an extreme edge.
2. There is a lot of data being generated in a large geographic area.
3. The information gathered needs to be analyzed, and it has to provoke a response in less than a second.

According to [77], an average cable provider in 2008 allowed 16 Mbps. Two years later, bandwidth increased to 31 Mbps. The last measurement was from 2019 with 325 Mbps. This behavior can be described by Nielsen's law, which says that each year, bandwidth increases by approximately 50%. Some authors even compare Nielsen's law with Moore's law, which describes computer power. Despite recent advances in bandwidth capacity, edge computing is still the most critical step for real-time cloud analysis. As shown in [78], a DDNN that involves the edge brings several benefits, such as fault tolerance, privacy, and reduced communication costs.

All and all, we conclude that edge and fog computing are the best ways to implement a NN in the cloud for applications that require fast analysis like AVs. Despite requiring more powerful end devices, or edge servers, these techniques reduce network traffic, allowing faster analysis of information, more privacy, and fewer transmission errors.

4.2. Parallelization of Neural Networks

The review of deep learning NNs presented in Section 2 showed that the accuracy of models could be improved by increasing the number of parameters and the scale (i.e. nodes, layers, etc.). However, these actions also derive slower training times, which is inefficient for adaptive autonomous driving scenarios. Therefore, we have acknowledged that deep learning models need to be parallelized to accelerate training and deployment. Indeed, deep learning models can be trained and executed in multiple GPUs. The key consideration to take advantage of their parallelization is to know how to divide the GPUs' tasks. Hence, three approaches should be considered if we want to train parallelized models: data parallelism, model parallelism, and data-model parallelism [79].

- Data parallelism: Data parallelism is a different kind of parallelism that, instead of relying on the process or task concurrency, is related to both the flow and the information structure. Each GPU uses the same model to train on different data subset. If there are too many computational nodes, it is necessary to reduce the learning rate to keep a smooth training process.
- Model parallelism: In model parallelism, each computational node is responsible for parts of the model by training the same data samples. The computational nodes communicate between them when the input of a neuron is from another computational node's output. However, the performance is worse than data parallelism, and the performance of the network will be decreased if it has too many nodes.
- Data-model parallelism: Both previous models have disadvantages, but they also have positive characteristics. Model parallelism could get good performance with many neuron activities and data parallelism is efficient with many weights.

In summary, data-model parallelism combines the best of the first two approaches and is more suited for AV tasks as long as we have a clear idea of where to apply each step. For instance, the convolutional layer of a CNN could use data parallelism (since it contains about 90% of the computation and 10% of the parameters). In contrast, the fully connected layer could only rely on model parallelism, because it contains 90% of the parameters and 10% of the computation.

4.3. Parallelization of the whole system

One thing that is rarely discussed in research is the likelihood of the proposed solution to be applicable in real-life settings. Factors such as the time required for the whole data to be processed, and if that time is good enough for real-time applications, are commonly discussed for a system to be deployed in the real world. However, hardware requirements for the processing are acceptable for these applications, and code parallelization accelerates the clocks in the processor at just the right rate speed (but not too much so that hardware operates at dangerous temperatures) need to be acknowledged as well. In this regard, thermal throttling can be considered as a viable solution. Thermal throttling is a safety protocol that the processor starts once the system reaches a determined temperature. It starts to divert power from processing the data towards the cooling system techniques to avoid any hardware malfunctions.

Although this technique is beneficial for the hardware, it derives from slower data processing (almost half of its capabilities). For examples, using thermal throttling in image processing might drop the frame rate from 30 FPS to 15 FPS (or less); this quality drop in autonomous driving could even be dangerous as it could result in lower recognition rates and maximum risk of accidents. This issue can be addressed using proper parallelization techniques in CPU and GPU architectures. It is implemented for data to be transferred smoothly throughout the system without frame drops and appropriate speed for real-time applications without compromising hardware's integrity. For instance, CUDA libraries from NVIDIA⁹ have excellent parallelisation frameworks to segment complex instructions into smaller ones; all of them are run concurrently in different GPU CUDA cores, giving the system the robustness and speed, it requires. From Google, TensorFlow API also has adequate parallelization protocols for NN training and deployment in GPU architectures. It can be used in the object detection, instance segmentation and RNN prediction tasks. The robotics operating system provides parallelization options in CPU for the SLAM, path planning,

⁹ NVIDIA, "CUDA Toolkit 10.1 original Archive", *NVIDIA developer*, [CUDA Toolkit 10.1 original Archive | NVIDIA Developer](https://developer.nvidia.com/cuda-toolkit-10.1-original-archive)

communication protocols and wireless data transfer for multiple path planning operations in the Frenét coordinate system.

4.4 First thoughts towards an approach for LSTM-based path planning prediction

After considering the conditions mentioned above, we propose a novel LSTM approach to predict both lateral displacement and linear velocity in a Frenét coordinate system. Many techniques have been proposed to achieve this (as described in Sections 2.10 and 3). Still, they are considered isolated ways to solve the problem as they do not consider integrating the overall pipeline. The reader must recall that vast amounts of data are required to perform this task and acknowledge that a critical concept is often ignored in this phase, i.e. causality. It is essential to consider that there is a high chance of a reckless driver making a poor judgment that could affect the AV or that the predictions of a single car won't change because another car changes the whole driving scenario to poor or unexpected choices. Even deployed systems, such as the Ego Car (described in Section 3) would be part of this causality scenario. Therefore, using the output from the Frenét coordinate system that was taken from object classification via YOLO and boundary classification with instance segmentation, a bidirectional LSTM can be fed with the Frenét data, both for lateral displacements and linear velocities.

These are four different LSTM models that feed their weights to each other and share weights between lateral displacements and linear velocities. As seen in Figure 18, the LSTM architecture is represented graphically.

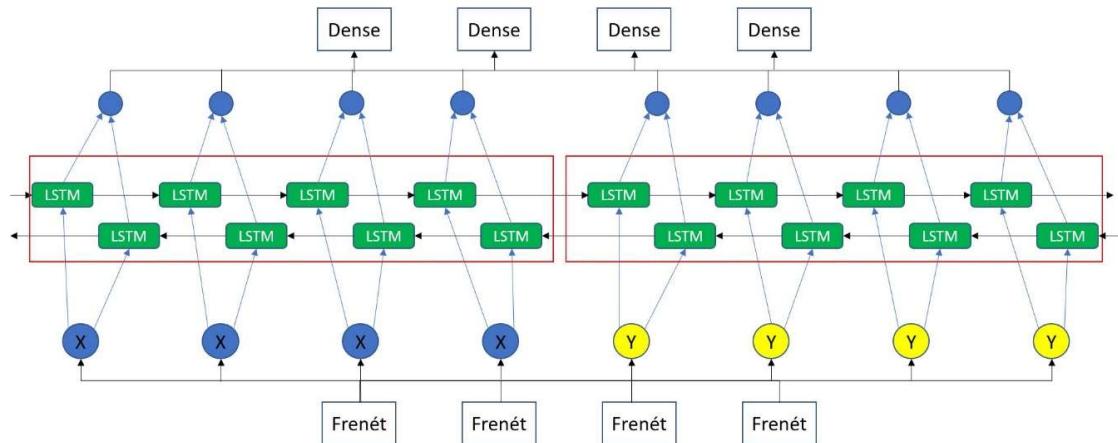


Figure 18. LSTM representation for the prediction task.

Each and every Frenét point of each and every car is going to be fed to the model, including the Ego Car points, and thus the prediction will take place considering every variable in the scenario, and changing its predictions if the causality demands it. This will derive in a new Frenét prediction point for every neuron in the LSTM architecture, and so the position, orientation, and velocity of every car can be predicted. As a result, the cars volume can be included in every future time instance. This output data can be used for decision making in an autonomous car, reducing enormously the number of accidents since causality is part now of the overall algorithm deployment and execution.

5. Conclusions

This paper presented a comprehensive and thorough review towards achieving autonomous driving by considering concepts such as image processing, NNs, and other deep learning models. Moreover, we acknowledged two key concepts that must be addressed to deploy these systems in real life, i.e.

DS and parallelization. The review was structured towards establishing the proper ways to evaluate these systems and structure their use in pipeline systems that attempt to perform path planning (by means of Frénét coordinates) and identify multiple objects on the road analyze the behavior of the driver. Although the review's scope is extensive and demanded a considerable amount of time to sort and select the proper references, we believe that the result is a comprehensive guide for novel and experienced researchers who intend to develop systems autonomous systems in the near future.

Our efforts and research in this domain do not conclude here; we are still researching the latest trends in this area and incorporating them into our proposed system's design. Also, we are considering how the latest developments, such as the COVID-19 pandemic, will affect these systems' development. The team comprises members working in different parts of the world. Therefore, our investigation in DS and parallelization techniques is not only intended for the sole purpose of improving the performance of the system, but also to better sort the training and deployment tasks in different machines and servers.

6. Table of Abbreviations

ADS: Autonomous Driving System

AV: Autonomous Vehicle

AP: Average Precision

mAP: Mean Average Precision

BC: Behavioral Cloning

BEV: Bird's Eye View

CDBN: Convolutional Deep Belief Network

CLDA: Constrained Linear Discriminant Analysis

CNN: Convolutional Neural Network

COCO-ODC: Common Objects in Context Object Detection Challenge

CRNN: Convolutional Recurrent Neural Network

DAG-SVM: Directed Acyclic Graph-Support Vector Machine

DARPA: Defense Advanced Research Projects Agency

DBN: Deep Belief Network

DDNN: Distributed Deep Neural Network

DS: Distributed System

DOG: Difference of Gaussian

FCN: Fully Convolutional Network

FP: False Positive

FN: False Negative

GAN: Generative Adversarial Network

IARA: Intelligent Autonomous Robotics Automobile

IDM: Intelligent Driver Model

IoU: Intersection over Union

LiDAR: Light Detection and Ranging

LSTM: Long Short-Term Memory

MOT: Multi-Object Tracking

M³OT: 3D Multi-Object Tracking

MV3D: Multi-View 3D Object Detection Network for Autonomous Driving

MVX-Net: Multimodal Voxel for 3D Object Detection Network

NHTSA: National Highway Traffic Safety Administration

NN: Neural Network

R-CNN: Region-Based Convolutional Neural Network

RL: Reinforcement Learning

RNN: Recurrent Neural Network

RPN: Region Proposal Network

RoI: Region of Interest

SAE: Society of Autonomous Engineers

SIFT: Scale Invariant Feature Transform

SLAM: Simulation Localization And Mapping

SURF: Speeded-Up Robust Features

TN: True Negative

TP: True Positive

VOC: Visual Object Classes

Author Contributions: ECGM, VMVE, DCS, SFRC, MPH, AZC & JAGG investigated, compiled papers and wrote the manuscript. RBB and CFMG supervised all developments and revised the manuscript.

Funding: This research received no external funding

Acknowledgments: Authors thank to Centro de Investigación en Microsistemas y Biodiseño (cimb.com.mx) at Tecnológico de Monterrey Campus Ciudad de México for all the support in the research process.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Singh, S. (2015) Critical Reasons for Crashes Investigated in the National Motor Vehicle Crash Causation Survey. National Highway Traffic Safety Administration. U.S. Department of Transportation. National Highway Traffic Safety Administration.
2. Rangesh, A. & Trivedi, M. (2019) No blind Spots: Full-Surround Multi-Object Tracking for Autonomous Vehicles Using Cameras & LiDARs.
3. Huang, Y. & Chen, Y. (2020) Autonomous Driving with Deep Learning: A Survey of State-of-art Technologies.
4. Yurtsever, E., Lambert, J., Carballo, A. & Takeda, K. (2020) A survey of Autonomous Driving: Common Practices and Emerging Technologies.
5. Badue, C., Guidolini, R., Vivacqua, R., Azevedo, P., Brito, V., Forechi, A. & Ferreira, A. (2019) Self-Driving Cars: A survey.
6. Lowe, D. G. (1999) Object recognition from local scale-invariant features. Proceedings of the International Conference on Computer Vision, Volume 2, pp. 1150–1157.
7. Bay, H., Ess, A., Tuytelaars, T. & Van Gool, L. (2006) Speeded Up Robust Features. Lecture Notes in Artificial Intelligence, Volume 3951, pp. 401-417.
8. Cortes Gallardo-Medina, E., Moreno-Garcia, C. F., Zhu, A., Chípuli-Silva, D., González-González, J. A., Morales-Ortiz, D., Fernández, S., Urriza, B., Valverde-López, J., Marín, A., et al. (2019) A Comparison of Feature Extractors for Panorama Stitching in an Autonomous Car Architecture. IEEE International Conference on Mechatronics, Electronics and Automotive Engineering (ICMEAE), Cuernavaca, Mexico.
9. Varghese, J. et al. (2015) Overview of Autonomous Vehicle Sensors and Systems. International Conference on Operations Excellence and Service Engineering.
10. Mazzari, V. (2019) What is LiDAR technology? Generation Robotics. Retrieved from: <https://blog.generationrobots.com/en/what-is-lidar-technology/>
11. Beltran J, et al. (2018) BirdNet: A 3D Object Detection Framework from LiDAR Information. IEEE International Conference on Intelligent Transportation Systems.
12. Chen, X. (2017) Multi-View 3D Object Detection Network for Autonomous Driving. IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
13. Tan, D. (2020) Camera-Lidar Projection: Navigating between 2D and 3D. Medium. Retrieved from: <https://medium.com/swlh/camera-lidar-projection-navigating-between-2d-and-3d-911c78167a94>
14. Starr, W. & Lattimer, B. Y. (2012) A Comparison of IR Stereo Vision and LiDAR for Use in Fire Environments.
15. Viitaniemi, V. & Laaksonen, J. (2008) Techniques for Image Classification, Object Detection and Object Segmentation.
16. Leonard, J. K. (2019). Image Classification and Object Detection Algorithm Based on Convolutional Neural Network. Science Insights, Volume 31, Issue 1, pp. 85–100.
17. Chen, C., Qin, C., Qiu, H., Tarroni, G., Duan, J., Bai, W., & Rueckert, D. (2020) Deep Learning for Cardiac Image Segmentation: A Review. Frontiers in Cardiovascular Medicine, Volume 7, Issue 25.
18. Tabian, I., Khodaei, Z. S., & Fu, H. (2019). A Convolutional Neural Network for Impact Detection and Characterization of Complex Composite Structures. Sensors, Volume 19, Issue 22, 4933.
19. Sabzekar, M., Ghasemigol, M., Naghibzadeh M. & Yazdi, H. S. (2013) Improved DAG-SVM: A New Method for Multiclass SVM Classification. ICAI.
20. Qian, D. (2005) Unsupervised Real Time Constrained Linear Discriminate Analysis to Hyper Spectral Image Classification. Pattern Reorganization, pp. 361-368.
21. Lecun, Y., Bottou, L. & Bengio, Y., (1998) Gradient-based Learning Applied to Document Recognition. Proceedings of IEEE, Volume 86, Issue 11, pp. 2278-2324.

22. Lee, H., Grosse, R., Ranganath, R., et al. (2009) Convolutional Deep Belief Networks for Scalable Unsupervised Learning of Hierarchical Representations. Proceedings of the 26th Annual International Conference on Machine Learning, pp. 609- 616.
23. Girshick, R., Donahue, J., Darrell, T., et al. (2014) Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. IEEE Conference on Computer Vision and Pattern Recognition, pp. 580-587.
24. Sheng, F. (2020) Review of LeNet-5: How to Design the Architecture Of CNN. Retrieved from: <https://towardsdatascience.com/review-of-lenet-5-how-to-design-the-architecture-of-cnn-8ee92ff760ac>
25. He, K., Gkioxari, G., Dollar, P. & Girshick, R. (2018) Mask R-CNN. Facebook AI Research, pp. 1-12, FAIR: IEEE.
26. Redmon, J., Divvala, S., Girshick, R. & Farhadi, A. (2015) You Only Look Once: Unified, Real-Time Object Detection.
27. Simonyan, J.K. & Zisserman, A. (2015) Very Deep Convolutional Networks for Large-Scale Image Recognition. International Conference on Learning Representations.
28. Pascanu, R., Mikolov, T. & Bengio, Y. (1998) On the Difficulty of Training Recurrent Neural Networks. *Chemin de la Tour*, pp. 1-16.
29. Mittal, A. (2019) Understanding RNN and LSTM. Retrieved from: <https://towardsdatascience.com/understanding-rnn-and-lstm-f7cdf6dfc14e>
30. Choi, K., Fazekas, G., Sandler, M., & Cho, K. (2017) Convolutional Recurrent Neural Networks for Music Classification. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 2392-2396.
31. Çakır, E., Parascandolo, G., Heittola, T., Huttunen, H. & Virtanen, T. (2017). Convolutional Recurrent Neural Networks for Polyphonic Sound Event Detection. IEEE/ACM Transactions on Audio, Speech, and Language Processing, Volume 25, Issue 6, pp. 1291-1303.
32. Maddula, R., Stivers, J., Mousavi, M., Ravindran, S., & Sa, V. D. (2017) Deep Recurrent Convolutional Neural Networks for Classifying P300 BCI signals. GBCIC. Retrieved from: <https://d3i71xaburhd42.cloudfront.net/8de0d05fea18ca27e53e28b94b8e31df4577f5d3/3-Figure2-1.png>
33. Zuo, Z. et al. (2015) Convolutional Recurrent Neural Networks: Learning Spatial Dependencies for Image Representation. IEEE Conference on Computer Vision and Pattern Recognition, pp. 18-26.
34. Hu, Z. et al. (2019) A CRNN Module for Hand Pose Estimation. Neurocomputing, Volume 333, pp. 157-168.
35. Olah, C. (2015) Understanding LSTM Networks. Retrieved from: <http://colah.github.io/posts/2015-08-Understanding-LSTMs/>
36. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014) Generative Adversarial Nets. Advances in Neural Information Processing Systems.
37. Mirza, M. & Osindero, S. (2014) Conditional Generative Adversarial Nets. arXiv.
38. Odena, A. (2016) Semi-supervised Learning with Generative Adversarial Networks. arXiv.
39. Odena, A., Olah, C. & Shlens, J. (2017) Conditional Image Synthesis with Auxiliary Classifier GANs. International Conference on Machine Learning.
40. Ali-Gombe, A., Elyan, E., Savoye, Y. & Jayne, C. (2018) Few-shot classifier GAN. International Joint Conference on Neural Networks (IJCNN).
41. Kulpatti, S. (2019). A Brief Introduction to GANs. Retireved form Sigmoid: <https://medium.com/sigmoid/a-brief-introduction-to-gans-and-how-to-code-them-2620ee465c30>
42. Yong, Fu., Tanfeng, Sun, Xinghao, Jiang, Ke, Xu, Paisong, He. (2018) Robust GAN-Face Detection Based on Dual-Channel CNN Network: <https://ieeexplore.ieee.org/abstract/document/8965991>
43. Ali-Gombe, A., Elyan, E. & Jayne, C. (2017) Fish Classification in Context of Noisy Images. Engineering Applications of Neural Networks, Volume CCIS 744, pp. 216-226.
44. Elyan, E., Jamieson, L. & Ali-Gombe, A. (2020) Deep Learning for Symbols Detection and Classification in Engineering Drawings. Neural Networks, Volume 129, pp. 91-102.
45. Coulouris, G., Dollimore, J., Kindberg, T., & Blair, G. (2012). *Distributed Systems Concepts and Design*.
46. Usman, M., Anjum, A., Farid, M. & Antonopoulos, N. (2018) Cloud-Based Video Analytics Using Convolutional Neural Networks. Cloud-based Video Analytics System.
47. Teerapittayanon, S., McDanel, B. & Kung, H. T. (2017) Distributed Deep Neural Networks over the Cloud, the Edge and End Devices.

48. Khandelwa, R. (2019) COCO and Pascal VOC data format for Object detection. *Towards Science*.
49. Zeng, N. (2018) An Introduction to Evaluation Metrics for Object Detection.
50. El Aidouni, M. (2019) Evaluating Object Detection Models: Guide to Performance Metrics. *Github*.
51. Wrling, W. et al. (2010) Optimal Trajectory Generation for Dynamic Street Scenarios in a Frenét Frame. *FAIR: IEEE*.
52. Zhou, Y. & Tuzel, O. (2017) VoxelNet End-to-End Learning for Point Cloud Based 3D Object Detection. *arXiv*.
53. Zhou, Y. & Tuzel, O. (2019) MVX-Net: Multimodal VoxelNet for 3D Object Detection.
54. Liu, T., Liao, Q., Gan, L., Ma, F., Cheng, J., Xie, W. et al. (2020) Hercules: An Autonomous Logistic Vehicle for Contact-less Goods Transportation During the Covid-19 Outbreak. *arXiv*.
55. Moraes, G., et al., "Image-Based Real-Time Path Generation Using Deep Neural Networks," 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, United Kingdom, 2020.
56. Cardoso, V. B. et al., "A Large-Scale Mapping Method Based on Deep Neural Networks Applied to Self-Driving Car Localization," 2020 International Joint Conference on Neural Networks (IJCNN), Glasgow, United Kingdom, 2020.
57. Weaver, C. (2020) Self-Driving Cars Learn to Read the Body Language of People on the Street. *IEEE SPECTRUM*.
58. Aranzeta-Ojeda L., Moreno-García, C.F., Granados-Reyes, A. & Bustamante-Bello R. (2011) Design, Development and Testing of a Low-Cost, High Sensitivity System for Neurodegenerative Disease Detection and Characterization. *International Conference on Microtechnologies and Medical Biology (MMB)* pp. 64–65.
59. Bustamante-Bello, R., Aranzeta-Ojeda, L. & Moreno-Garcia, C.F. (2011) Design and Development of a Low-Cost, High Sensitivity Device for Neurodegenerative Disease Detection. *24th IEEE International Conference Micro Electro Mechanical Systems (MEMS)*.
60. Barea R., Bergasa, L. M., Romera, E., López-Guillén, E., Perez, O., Tradacete, M. & López, J. (2019) Integrating State-of-the-Art CNNs for Multi-Sensor 3D Vehicle Detection in Real Autonomous Driving Environments. *IEEE Intelligent Transportation Systems Conference (ITSC)*, Volume 1, Issue 7.
61. Shah, N., Shankar, A. & Jae-hong Park. (2020) Detecting Drivable Area for Autonomous Vehicles. *Duke University Department of ECE*, *arXiv*.
62. Meng, X., Lee, K. K., & Xu, Y. (2006). Human Driving Behavior Recognition Based on Hidden Markov Models. *2006 IEEE International Conference on Robotics and Biomimetics*.
63. Curiel-Ramirez, L.A., et al. (2020) End-to-End Automated Guided Modular Vehicle. *Applied Sciences*, *MDPI*.
64. Sutton, R.S. & Barto, A. G. (2018) Reinforcement Learning: An introduction. Second Edition. Cambridge, MA: The MIT Press.
65. Hoel, C. J., Wolff, K., & Laine, L. (2018). Automated Speed and Lane Change Decision Making using Deep Reinforcement Learning. *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*.
66. Sallab, A. E., Abdou, M., Perot, E., & Yogamani, S. (2017) Deep Reinforcement Learning Framework for Autonomous Driving. *Electronic Imaging*, vol. 2017, no. 19, pp. 70–76.
67. Bai, Z. & Cai, B. (2019) Deep Learning Based Motion Planning for Autonomous Vehicle Using Spatiotemporal LSTM Network. *arXiv*.
68. Atlché, F. & de la Fortelle, A. (2018) An LSTM Network for Highway Trajectory Prediction. *arXiv*.
69. Yan, S., Teng, Y., Smith, J. & Zhang, B. (2016) Driver behavior recognition based on deep convolutional neural networks. *12th International Conference on Natural Computation, Fuzzy Systems and Knowledge Discovery (ICNC-FSKD)*, pp. 636-641.
70. Naqvi, R. A., Arsalan, M., Rehman, A., Paul, A., Rehman, A. U., & Loh, W.-K. (n.d.). Deep learning-based drivers emotion classification system in time series data for remote applications. *Remote Sensing*, 12(3). <https://doi.org.biblioteca-ils.tec.mx/10.3390/rs12030587>
71. Hu, Y., Lu, M., & Lu, X. (2020). Feature refinement for image-based driver action recognition via multi-scale attention convolutional neural network. *Signal Processing Image Communications*, issue 81.
72. Kim, C. M., Hong, E. J., Chung, K., & Park, R. C. (n.d.). Driver facial expression analysis using LFA-CRNN-based feature extraction for health-risk decisions. *Applied Sciences*, volume 10, issue 8.

73. Affectiva (n.d) Affectiva Automotive AI for Driver Monitoring Systems. Retrieved from: <https://www.affectiva.com/product/affectiva-automotive-ai-for-driver-monitoring-solutions/>
74. Liu, S., Tang, J., Wang, C., Wang, Q., and Gaudiot, J.L. (2017) "Implementing a Cloud Platform for Autonomous Driving". Cornell University. Retrieved from <https://arxiv.org/abs/1704.02696>
75. Kumar, S., Gollakota, S. and Katabi, D. (2012). "A Cloud-Assisted Design for Autonomous Driving". Massachusetts Institute of Technology. Retrieved from https://www.researchgate.net/publication/254464262_A_cloud-assisted_design_for_autonomous_driving
76. Cisco. (2015). Fog Computing and the Internet of Things: Extend the Cloud to Where the Things Are. White Paper. Retrieved from: https://www.cisco.com/c/dam/en_us/solutions/trends/iot/docs/computing-overview.pdf
77. Nielsen, J. (2019). Nielsen's Law of Internet Bandwidth. Nielsen Norman Group. Retrieved from <https://www.nngroup.com/articles/law-of-bandwidth/>
78. Teerapittayanon, S., McDanel, B. & Kung, H. T. (2017) Distributed Deep Neural Networks over the Cloud, the Edge and End Devices. IEEE 37th International Conference on Distributed Computing Systems (ICDCS).
79. Li, X., Zhang, G., Li, W. et al. (2016). Big Data: Chapter 4 - Deep Learning and its Parallelization.
80. Wikipedia, The Free Encyclopedia; Frenet-Serret formulas, Definitions.