
An Intelligent Modular Real-Time Vision-Based System for Environment Perception

Amirhossein Kazerouni*

Iran University of Science and Technology
amirhossein477@gmail.com

Amirhossein Heydarian*

Amirkabir University of Technology
ah.heydarian@aut.ac.ir

Milad Soltany*

Iran University of Science and Technology
soltany.m.99@gmail.com

Aida Mohammadshahi*

University of Calgary
aida.mohammadshahi@ucalgary.ca

Abbas Omid*

University of Calgary
abbas.omidi@ucalgary.ca

Saeed Ebadollahi

Iran University of Science and Technology
s_ebadollahi@iust.ac.ir

Abstract

A significant portion of driving hazards is caused by human error and disregard for local driving regulations; Consequently, an intelligent assistance system can be beneficial. This paper proposes a novel vision-based modular package to ensure drivers' safety by perceiving the environment. Each module is designed based on accuracy and inference time to deliver real-time performance. As a result, the proposed system can be implemented on a wide range of vehicles with minimum hardware requirements. Our modular package comprises four main sections: lane detection, object detection, segmentation, and monocular depth estimation. Each section is accompanied by novel techniques to improve the accuracy of others along with the entire system. Furthermore, a GUI is developed to display perceived information to the driver. In addition to using public datasets, like BDD100K, we have also collected and annotated a local dataset that we utilize to fine-tune and evaluate our system. We show that the accuracy of our system is above 80% in all the sections. Our code and data are available on [GitHub](#).

1 Introduction

As the number of vehicles on the road has grown in recent years, traffic violations, accidents, and fatalities have increased considerably [13]. However, along with the growth in urban traffic, human error plays a significant role in the frequency of road casualties and offenses, in a way that careless driving and disregard for traffic signs account for more than 70% of street accidents [6]. Hence, providing a realistic solution to improve driving accuracy and road safety could be highly beneficial. Artificial intelligence has advanced many fields, including the automotive industry [2, 22, 23]. Some companies have offered autonomous vehicles as an alternative to human driving [24], but they are still not commonly used by the public due to their high cost.

Road lane detection is a significant part of the perception system to determine the vehicle's position and its desired trajectory on the road, which remains problematic in intelligent vehicles. The challenges of this problem include disordered, fading, and unreasonable road markings, along with various lighting conditions. Furthermore, road lines might be occluded by other cars and obstacles. Early lane detection methods include classical image processing techniques based on image edges, and the Hough transform [32, 38, 42, 59, 54]. Despite their simplicity, these algorithms require environment-dependent hyperparameter adjustments. [41] presents a hybrid method based on object

*Equal contribution

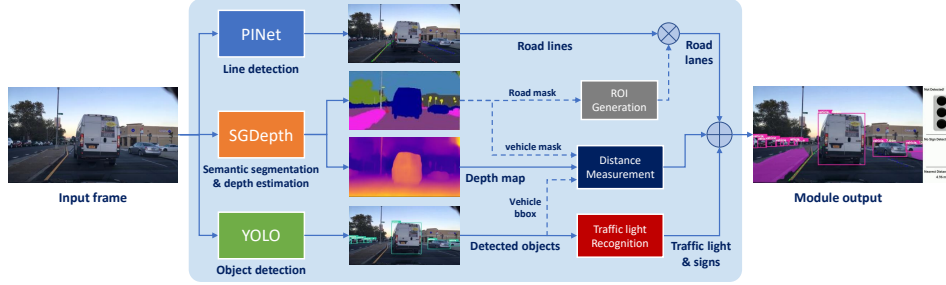


Figure 1: **Overview of the proposed system.** Continuous lines represent system modules’ main data path, and dashed lines denote the additional data, which aids in the accuracy of the system.

detection to improve lane detection accuracy. On the other hand, deep learning-based methods such as Point Instance Network (PINet) [31] and LaneATT [55] offer precise and robust performance in multiple scenarios with different lighting conditions, and occlusions [44, 36, 1, 20]. As a result, deep-learning-based approaches are considered a better choice for practical applications.

Object Detection methods are used to detect various objects in the environment. Most detectors feed an image into an artificial neural network and output bounding box locations for each object present in the image and probability scores corresponding to each object class. The detector’s head connecting to the backbone is usually one or two-stage. One-stage detectors, for instance, YOLO and SSD, are generally much faster than their two-stage counterparts [26, 7, 47], while two-stage detectors, like the R-CNN family, tend to yield more outstanding accuracy scores [37, 34, 16, 15, 50]. Our first use of object detection is detecting pedestrians and other vehicles, static and dynamic; this is done to avoid collisions with other objects on the road. Secondly, traffic indicators are detected to ensure vehicles follow the regulations on different roads [12].

An intelligent system should accurately identify the road-sidewalk boundaries and measure its distance from detected objects to prevent accidents. In this case, pixel-wise classification of an image (semantic segmentation) offers impressive accuracy in determining the boundary by recognizing the sidewalk [40, 8, 53, 61, 51, 5]. In addition, perspective transform approaches [28, 57] that estimate the distance based on the warped image’s pixel spacing and monocular depth estimation methods [58, 39, 10, 33] that measure the distance based on the depth map and camera parameters have shown promising results. However, the accuracy of perspective-transform-based approaches is highly dependent on the situation and cannot perform well in all conditions. Moreover, it is not a computationally cost-effective solution to measure the distance and detect the sidewalk individually through different models. For this purpose, models based on multi-task learning can be advantageous since they output both semantic segmentation and depth estimation in one model [64, 62, 9]. These models simultaneously enhance the accuracy of segmentation and depth estimation through their mutual effect during training. For instance, SGDepth [30] outputs the estimated depth map and semantic segmentation from a single input image. In this research, we have developed a comprehensive package for real-time environmental perception based on computer vision techniques to assist drivers in minimizing driving faults and violations. To the best of our knowledge, we have provided a novel combination of traits, and only a few studies are available on the permutations of features mentioned in this paper. For instance, compared to the YOLOP[60] network, this network lacks depth estimation and distance measurement, whereas our package is more extensive. There are two aspects to our real-time intelligent package: software and hardware. The software section consists of four main phases. For the road lane detection part, our module benefits from the PINet [31] neural network, which is trained on the CULane [52] dataset and fine-tuned on our collected local dataset. For the object detection section, we utilize YOLOv5 [27] for detecting vehicles, pedestrians, and traffic signs. In the third phase, we use the SGDepth [30] network for segmentation tasks and recognizing the sidewalks. Finally, we introduce a novel approach for measuring the distance to the surrounding cars based on the monocular depth estimation output from the SGDepth. In terms of hardware, this module uses only one camera and a mid-range GPU, reducing costs and making it effortless to implement on various vehicles.

2 Background

2.1 PINet

The Point Instance Network (PINet) detects traffic lines regardless of their number [31]. It generates points on lanes and separates them into distinct instances. As illustrated in Fig 2, three output




Figure 2: The overview of PINet [31].




Figure 3: The overview of SGDepth [30].

branches are included in this network: a confidence branch, an offset branch, and an embedding branch. Predicting the exact points of traffic lines is what the confidence and offset branches do. The embedding branch creates the embedding features of the predicted points, which are provided in the clustering process to differentiate each instance. First, the input RGB image enters the resizing network, and the sequence of three convolution layers compresses the image into a lower size. After each convolution layer, PReLU [21] and Batch Normalization [25] are used. Then, the predicting network receives the resizing network output. It predicts the exact points on the traffic lines as well as embedding features. There are multiple hourglass modules in this network, each with an encoder, decoder, three output branches, and some skip-connections. The predicting network can include any number of hourglass modules, and all of them are trained concurrently by the same loss function. Therefore, in the case of running the trained model on a system with limited computing power, the network can be cut and transferred without extra training. Hourglass blocks have three types of bottlenecks: same, down, and up bottlenecks. The output of the same bottleneck is the same size as the input. In the encoder, the down bottleneck is used for down-sampling, with the first layer of a convolution layer, and in the up-sampling layer, a transposed convolution is used for the up bottleneck. Each output branch has its own channel (confidence: 1, offset: 2, embedding: 4). The associated loss function is applied based on the output branch’s goal, and each confidence output is passed on to the following block. The PINet network loss function equation consists of four different loss functions. Three of them (confidence, offset, and feature loss functions) are parallel and are applied to the network output branch. The other one (distillation loss function) optimizes the knowledge-learning process from the deepest hourglass, which is beneficial if the prediction network is cut to a lighter one. The total loss function equals the weighted sum of the above four loss functions, as shown below:

$$L_{total} = \gamma_e L_{exist} + \gamma_n L_{non-exist} + \gamma_o L_{offset} + \gamma_f L_{feature} + \gamma_d L_{distillation}. \quad (1)$$

Where constant coefficients are obtained experimentally.

2.2 YOLO

YOLOv5 [27] network is a single-stage detector which consists of three vital parts. The backbone, the initial element, extracts image features at various scales. The neck, which merges the retrieved features, is the second element. And the network head predicts the box and class of each object in the picture using the features coming from the neck. Eventually, Yolo encodes all the information into a single tensor for each image [48]. YOLO models the detection problem as regression by dividing the input image into an $S \times S$ grid. Each grid cell predicts B bounding boxes with x , y , w , and h , and an “objectness” score $P(Object)$, which indicates whether or not the grid cell contains an object. In addition, a conditional probability $P(Class | Object)$ is predicted for the class of the object associated with each grid cell. Therefore, YOLO outputs $B \times 5$ parameters and C class probabilities for each grid cell. These predictions are represented by a tensor with the size $S \times S \times (B \times 5 + C)$. Non-Maximum Suppression (NMS) and thresholding are the last steps in generating final object detection predictions [48]. The network parameters are trained by minimizing a three-element loss function called GIoU, Objectness, and Classification. The GIoU element minimizes bounding box prediction error, and the objectness element determines the existence of an object in a grid cell. Finally, the classification error is responsible for object class prediction. YOLOv5 benefits from previous version features, but it is implemented in PyTorch rather than Darknet, making it more flexible with computer vision libraries.

2.3 SGDepth

SGDepth is a novel monocular self-supervised approach to estimate depth information from single images. This process is semantically-guided, meaning it utilizes information obtained from segmentation maps and does not require depth labels. As shown in Fig. 3, the approach consists of two main

components: the depth part, which is trained in a self-supervised fashion, and the segmentation part, which utilizes a supervised training scheme. For models to estimate depth maps from sequences of images, the world must be static, i.e., in two consecutive frames of a video, all objects must remain in their absolute positions. Moving dynamic-class (DC) entities, including passing vehicles and pedestrians, violate the static world assumptions. Hence, while being necessary for self-supervised depth estimation, correct projections could not be calculated between sequences of frames. In the training phase, SGDepth ignores such objects in their optimization.

Self-Supervised Monocular Depth Estimation. Instead of training the model on depth labels in a supervised manner, the predicted depth maps are then considered as geometric properties of the environment to warp the preceding and succeeding frames $\mathbf{x}_{t'}$, with $t' \in \mathcal{T}' = \{t-1, t+1\}$ to \mathbf{x}_t at time t . Afterward, a photometric loss J_t^{ph} is computed between \mathbf{x}_t and $\mathbf{x}_{t' \rightarrow t}$; this is to ensure that the transformed images $\mathbf{x}_{t' \rightarrow t}$, are as close as possible to \mathbf{x}_t . On top of that, a smoothness loss is responsible for making sure that nearby pixels have similar depth values [17, 18].

Supervised Semantic Segmentation. A segmentation mask $\mathbf{m}_t \in \mathcal{S}^{H \times W}$, where \mathcal{S} is a set of classes, is obtained by assigning a class to each pixel coordinate. This is done through a supervised training method, and the segmentation head (Fig. 3) outputs $\mathbf{y}_t \in \mathbb{I}^{H \times W \times S}$ are compared to the ground truth labels $\bar{\mathbf{y}}_t$ using a weighted cross-entropy loss [43].

Semantic Guidance. As shown in Fig. 3, there are two decoders attached to one encoder (Feature Extractor block). In the backward propagation stage, the gradients that return from the two decoders are scaled to form $\mathbf{g}^{\text{total}}$ that is propagated back into the decoder. This is how multi-task training is done across the two domains. To deal with moving DC objects, projected segmentation maps are calculated using nearest-neighbor sampling. In addition, a DC object mask $\boldsymbol{\mu}_t \in \{0, 1\}^{H \times W}$ is computed that has zero values for all pixel coordinates that belong to a DC class \mathcal{S}_{DC} in either one of the three frames. This mask is then applied to compute a semantically-masked photometric loss. If a DC object has moved in two consecutive frames, the warped semantic $m_{t' \rightarrow t, i}$ and the semantic mask m_t will be inconsistent. To measure this, intersection over union (IoU) of DC objects in $m_{t' \rightarrow t, i}$ and m_t is calculated and denoted as $\Lambda_{t, t'}$. Then a threshold $\theta_\Lambda \in [0, 1]$ is used to decide whether a frame is static or dynamic, where photometric loss or masked photometric loss is applied, respectively. After this, the total loss will be a combination of the cross-entropy loss and smoothness loss, as well as the photometric losses.

3 Method

As previously mentioned, our modular intelligent system includes four parts. In the lane detection section, a series of operations are carried out to display the road lines on the module screen. This section aims to help drivers keep their position in the lane and notice their deviation from the off-road side. In the object detection section, surrounding objects, such as obstacles on the road and traffic signs, are detected so that the system alerts drivers to route appropriately and follow traffic laws. In the segmentation section, the segments of the sidewalks are shown on the screen during the semantic segmentation process, which aids drivers in precepting their environment. Finally, in the distance measurement part, the distance from the nearby cars is estimated using a monocular depth estimation approach to avoid collisions and accidents.

Each of the four module’s parts utilizes a base deep learning model, but some do not have a robust and proper result, necessitating some creativity to solve problems. The method section describes how issues are overcome and what novelties are employed to boost the models’ performance and create an intelligent modular system. Fig. 1 illustrates the overall system diagram.

3.1 Lane Detection

We have used the PINet model for the lane detection part because, as observed from the experiments done and given in Table 5, it has demonstrated greater accuracy than other rivals while maintaining the required Frame Per Second (FPS). However, the model requires considerable processing power, which results in a noticeable drop in FPS in the performance of our entire module, so PINet is considered a trade-off between accuracy and FPS.

Although PINet shows accurate results for determining road lines, some environmental conditions, like poor illumination, cause severe errors in PINet prediction. In these cases, traffic lines are detected outside the road space. As a result, a novel dynamic region of interest (ROI) is designed and applied to the image to avoid off-road lines being identified. Our experiments show that all issues are addressed after adding this modified ROI to the PINet output, which can be seen in Fig. 4.

Building the modified ROI begins with deriving the road mask from the segmentation map output of the SGDepth. The road mask may have some empty areas caused by removing vehicles and objects from the road. Next, the convex hull approach is used on the road segment to obtain a convex space covering all the road mask edges [4]. After this, a modified and seamless ROI is obtained from the road by filling the convex space gaps. The utilized convex hull formula is as follows:

$$\text{ROI} = \text{Co}(X) = \left\{ \sum_{i=1}^q \alpha_i x_i \mid x_i \in X, \alpha_i \geq 0, \sum_{i=1}^q \alpha_i = 1 \right\} \quad (2)$$

Where $X = \{x_1, x_2, \dots, x_q\}$ denotes a set of points whose covered area is extracted by the convex hull. Morphological erosion and dilation are also applied for noise canceling and creating soft margins for the mask [19]. The erosion of the binary mask by the structuring element B is defined by:

$$\text{ROI} \ominus B = \{z \in E \mid B_z \subseteq \text{ROI}\} \quad (3)$$

Where E is a Euclidean space, the structuring element B is a circular disc in the plane, and B_z is the translation of B by the vector z . Also, the dilation of the binary mask by the structuring element B is defined by:

$$\text{ROI} \oplus B = \{Z \in E \mid (B^s)_z \cap \text{ROI} \neq \emptyset\} \quad (4)$$

Where E and B are the same as above, and B^s denotes the symmetric of B . Finally, by applying the modified ROI to the input image, the surroundings are removed, and only the in-road detected lines remain, which specify lanes.

Another challenge at this stage is facing roads and ROIs that do not contain traffic lines. To solve this, we introduce a new dataset containing local images collected by us in which roads do not include traffic lines. In the next step, traffic lines are labeled manually, and the network is fine-tuned on a combination of our introduced dataset with pre-existing datasets. As a result, due to the key point extraction nature of PINet, key points, even in the absence of white lines and in noisy conditions, are extracted well, and the fine-tuned PINet is able to divide the road into hypothetical lanes. The dataset introduced in this section is available on the project’s GitHub.

3.2 Object Detection

For the object detection part, the YOLO network has been used, which is a state-of-the-art, real-time object detection system. We have chosen the YOLOv5 version because, as shown in Table 6, it provides significantly higher accuracy and FPS compared to its counterparts. Utilizing the pre-trained YOLO object detection model, vehicles and pedestrians are detected, and bounding boxes are drawn around them on the monitor to alert the driver. However, identifying traffic signs is challenging since solely detecting them is not sufficient; it is necessary to recognize their meaning so that the driver can follow them. As a result, the YOLO network has been fine-tuned using a combination of multiple traffic signs datasets to both detect them and also identify the type of traffic signs.

The final acquired dataset includes fifteen traffic signs, most of which are similar in shape (e.g., circular with a red margin) but different in meaning. As a result of fine-tuning the model on this dataset, the traffic signs are accurately identified. After a sign is detected, the name of the sign is displayed on the driver’s monitor, depending on the type of detected sign.

Also, for identifying the traffic lights’ colors, after using the YOLO and detecting the traffic lights, we use a lightweight CNN-based classifier based on RegNetY002 [45] to classify four classes of red, yellow, green, and off states. When a traffic light is detected, the identified color of the light determines the displayed message. If the light is green, the word “pass” displays, and if the light is yellow or red, the words “warning” or “stop” display, respectively.

The results of the object detection section are improved compared to the network’s initial version after gathering the datasets and applying the modifications mentioned above, as shown in Table 2.

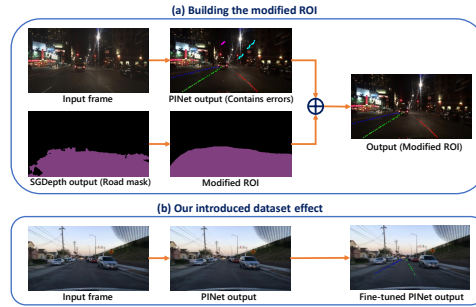


Figure 4: Summary of the implemented method in the lane detection section.

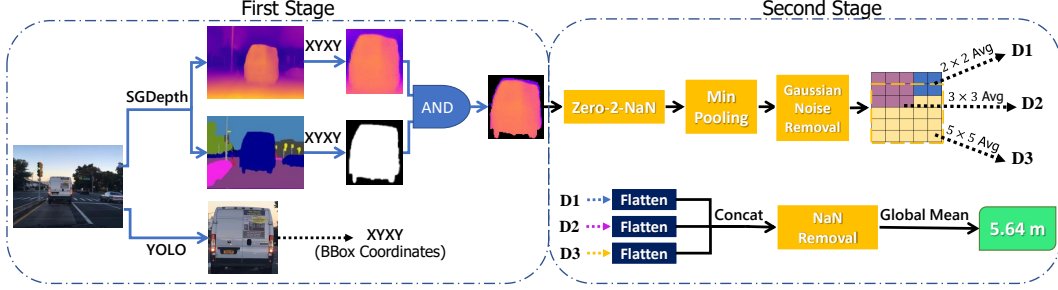


Figure 5: **The overview of distance measurement approach.** In the first stage, the object of interest is derived from the input image. Then the distance from the perceived object is determined in the second stage.

3.3 Segmentation

We utilize image segmentation for two primary purposes: a) avoiding curb collisions by detecting sidewalks and b) accurately measuring the distance from the vehicle to other objects. In the SGDepth model, which is a multi-task learning network, both the tasks of measuring the distance through the monocular depth estimation approach and separating the sidewalk segment from the road are done at the same time. Typically, environmental conditions have a significant influence on the accuracy of the model. Still, by employing this multi-tasking method, the extracted features of both tasks are shared in the network, leading to generalization and better results.

SGDepth is trained on the Cityscapes dataset [11], and we use the sidewalk, car, and bus classes to segment the scene with respect to said class labels. In addition, sidewalks will be colorized and displayed on the monitor. Since the segmentation is usually quite noisy, the approach still needs some novelties and modifications to perform appropriately. We deploy contour analysis and convex hulls to alleviate the sharp edges segmentation map. First, we use morphological functions to dilate the segmented area for each class and fill the holes inside the map that are the outcome of environmental noise. After that, the contours belonging to each semantic class are detected and sorted based on their encapsulating area, and contours whose area is smaller than a constant threshold are discarded. The remaining contours are then converted to convex hulls to compensate for the concavity that might occur in some instances. This way, the imperfections present in the initial segmentation map will be resolved considerably, and a much smoother segmentation map will be obtained.

3.4 Distance Measurement

In this section, we propose a novel hierarchical approach that exploits SGDepth and YOLO outputs to attain an accurate distance measurement. Our method consists of two major stages. The first stage is responsible for perceiving the object of interest in the scene. A more in-depth object analysis is achieved in the second stage to obtain the final distance measurement.

Since SGDepth operates on the entire scene, it is vital to identify the surrounding objects first. Therefore, as seen in Fig. 5, the object detection output is used to crop the detected object depth map and its segmentation mask based on the bounding box coordinates. We then take advantage of the acquired mask and apply it to the detected depth map to remove the background.

In order to eliminate the effect of zero-pixel values, we replace them with NaN to not affect the distance measurement calculations. The obtained mask may contain pixels from the background that are assumed to be outliers, and their presence may negatively affect distance measurement. To mitigate their impact, we apply min pooling on the depth map with a kernel size of 3×3 in the second stage. Min pooling operation with the kernel size of K is defined as follows

$$\hat{I}_{s,p} = \min(I_{i+K \times s, j+K \times p}) \quad \text{Min Pooling : } I^{H \times W} \rightarrow \hat{I}^{\lceil \frac{H}{K} \rceil \times \lceil \frac{W}{K} \rceil} \quad (5)$$

for $i = j = 0, 1, \dots, K - 1$, where $s \in S = \{0, 1, \dots, \lceil \frac{H}{K} \rceil - 1\}$, and $p \in P = \{0, 1, \dots, \lceil \frac{W}{K} \rceil - 1\}$ determine the set of new coordinates after applying min pooling. To ensure noisy depth points are not involved in the measurement procedure, we propose a Gaussian Noise Removal module that fits a Gaussian distribution on the depth values. Points whose distances are within an interval of $[\mu - 2\sigma, \mu + 2\sigma]$ are known as inliers and remain, while the rest of the points will be excluded and given NaN value to not participate in the distance measurement process. The goal of the Gaussian Noise Removal module is to find inlier points such that

$$\text{Inlier Points} = \left\{ \hat{I}_{s,p} \mid \mu - 2\sigma < \hat{I}_{s,p} < \mu + 2\sigma \right\} \quad (6)$$

Table 1: Evaluation of the sidewalk and lane segmentation sections.

		IoU	Confidence
Lane	Local	0.861	91.38%
	BDD100K	0.842	89.11%
Sidewalk	Local	0.795	85.34%
	BDD100K	0.762	80.23%

Table 2: Evaluation of the object detection section.

		Precision	Recall	F1-Score	Accuracy
Vehicles	Local	96.43%	85.71%	90.75%	83.08%
	BDD100K	94.92%	85.62%	90.03%	81.87%
Pedestrians	Local	95.24%	90.91%	93.02%	86.96%
	BDD100K	92.59%	89.28%	90.91%	83.33%
Traffic Signs	Local	93.33%	87.5%	90.32%	82.35%
	BDD100K	88.88%	94.12%	91.43%	84.21%
Traffic Lights	Local	94.74%	90.00%	92.31%	85.71%
	BDD100K	93.55%	85.29%	89.23%	80.56%

Table 3: Evaluation of the distance measurement section.

		mean RA	Accuracy
Distance	Local	84.30%	88.37%
	BDD100K	78.04%	82.91%

where σ and μ denote standard deviation and mean, respectively. In addition, it is necessary to apply average pooling to the acquired depth map for two main reasons, (i) to smooth the depth values and ignore the remaining sharp values, and (ii) to achieve an accurate value to report as distance. However, one of the challenging problems of estimating the distance is the object’s image size, and applying the prior min pooling would reduce its spatial dimensions by a factor of three. As a result, the kernel size of average pooling should be compatible with the object depth map size to guarantee we have enough depth points for the measurement step. Hence, to satisfy the objects with large and small depth map sizes, we propose a grouped average pooling with different kernel sizes, which are 2×2 , 3×3 , and 5×5 . The output of each average pooling is then flattened and concatenated. The NaN Removal module is subsequently employed to eradicate NaN values. Eventually, the distance to the surrounding objects is obtained by global averaging the resultant depth points.

4 Experiments

For traffic sign detection, we needed both richness and diversity in the signs. To this end, we used a combination of the DFG Traffic Sign Dataset [56] and Traffic-Sign Detection and Classification in the Wild dataset [66]. Furthermore, we used the Common Objects in Context (COCO) dataset [35] for the detection of pedestrians and other vehicles. To evaluate the overall performance of our proposed system in all sections, we combined the BDD100K dataset with our collected local dataset. BDD100K [63] is an extensive collection of 100,000 short videos of driving in various weather and driving conditions. Also, our comprehensive local dataset contains 100 fifteen-second 30 FPS videos of vehicles, roads, traffic signs, pedestrians, and sidewalks.

4.1 Hardware Configuration

The proposed modular package requires a high-resolution camera, a GPU, and a display. A 1080p camera captures images from the vehicle’s perspective, an Nvidia GTX 1660 Ti GPU is in charge of tensor processing in neural networks, and a display indicates the system’s results and any necessary alarms to the user. In real-world experiments, some additional hardware components are required for implementation, so a laptop is employed as the hardware framework in the vehicle to avoid this complexity. Fig. 6 shows the mentioned system implemented in a standard vehicle.

4.2 Experimental Results

Our intelligent system evaluation consists of three phases: assessment of area recognition sections, object detection parts, and distance measurement. Evaluations are performed on 500 frames extracted from the system’s output on the BDD100K video dataset and our local dataset. The area recognition section includes detecting sidewalks and traffic lanes. The object detection section comprises identifying pedestrians, parked and moving vehicles, traffic signs, and traffic lights with their colors. Finally, distance labels have been utilized to evaluate the distance measurement section. The assessment methods for each step and the result tables are explained below.

Lane and sidewalk segmentation. The segmentation of lane areas and sidewalk is evaluated using the Intersection over Union (IoU) metric, which is the intersection of the predicted and labeled masks per their union. The road lane masks are derived from the regions between the detected lines located in the modified ROI, and the sidewalk masks are the SGDepth modified outputs. These masks are




Figure 6: Implementation of the system in the urban environment. The hardware components include the camera, and the laptop, as a platform for running the package. The output information is displayed to the driver using the GUI.



Figure 7: The proposed system results on the BDD100K dataset.

Table 4: Real-time Evaluation of the System

Method	YOLOv5 (Object)	YOLOv5 (Sign)	PINet	SGDepth	FPS
Ours	✓	✓	✓	✓	7.56
A	✓	✗	✗	✗	66.06
B	✗	✓	✗	✗	69.40
C	✗	✗	✓	✗	23.58
D	✗	✗	✗	✓	47.48




Figure 8: The proposed system results on the local dataset.

Table 5: Evaluation of different lane detection approaches. NVIDIA Tesla T4 is used for evaluation.

		F1	IoU	Confidence	FPS
CLRNet [65]	<i>Local</i>	48.43%	32.05%	37.38%	47.55
	<i>BDD100K</i>	50.90%	33.76%	39.83%	
CondLaneNet [36]	<i>Local</i>	46.52%	32.52%	38.26%	128.75
	<i>BDD100K</i>	54.66%	37.04%	42.61%	
PINet [31]	<i>Local</i>	82.57%	74.14%	82.21%	24.67
	<i>BDD100K</i>	76.29%	75.84%	83.55%	

compared to their labels based on the IoU criterion, and a 0.5 threshold is utilized to calculate output accuracy. The predictions with IoU greater than 0.5 are regarded as valid, and with IoU less than 0.5 are considered false. The acquired results are displayed in Table 1.

Object detection. The performance of our system in object detection is evaluated using precision, recall, F1-score, and accuracy criteria. Notably, detected traffic lights are evaluated as a correct prediction only if their color is classified correctly, as recognizing traffic lights without interpreting their color would not result in getting the right traffic message. The findings are shown in Table 2, which indicates that the system’s accuracy is not less than 80% in any subsection.

Distance measurement. We utilize relative accuracy (RA) as a metric for the accuracy of the distance measurement section of our proposed system. Relative accuracy is defined using the following formula:

$$RA = 1 - \frac{|Distance_{actual} - Distance_{predicted}|}{Distance_{actual}} \quad (7)$$

If RA is greater than 0.8, then the predicted value for distance is considered correct. Afterward, the accuracy of our system is the percentage of the correct predictions in the entire dataset. The system accuracy and mean relative accuracy are reported for both local and BDD-100k datasets in Table 3.

Figs.7-8 show the performance results of the proposed perception system on local and BDD100K datasets. A graphical user interface (GUI) is also placed on the right side to represent traffic information to the driver. The GUI information includes traffic light status, traffic signs, and the distance to the nearest car so that the driver can be aware of the situation at a glance.

5 Ablation Study

Real-time Evaluation of the System. Table 4 presents the FPS of the whole system and each module individually. Notably, the FPS value for each module is computed from the time spent on preprocessing, model output, and postprocessing. According to Table 4, our system achieves a noticeable 7.56 average FPS over BDD100K and local datasets while demonstrating impressive quantitative and qualitative results. A considerable amount of processing time is spent by the PINet model that we have chosen due to its distinctive outputs. Although some other models operate in real-time, they do not have the generalization ability shown by PINet (see Table 5). In addition, part of the system’s processing is utilized to display the UI and model outputs, all of which reduce the processing time of the entire system. Nevertheless, the 7.51 average FPS is a reasonable quantity for running the model in real-time and obtaining satisfactory results. Furthermore, Human Reaction

Time (HRT) in driving is the time between when the driver is placed in a critical situation and when he decides to take action [29]. While a constant value for HRT cannot be reported, different methods report values between 1.27s to 1.55 [46, 3]. We believe that our system is real-time because its overall response time is more than six times less than that of humans. It is possible to connect our module to the vehicle’s braking system so that it can slow down in critical situations; this adds the benefit of much faster response times compared to that of humans, rendering our system real-time.

Ablation on Lane Detection We evaluate different SOTA approaches in lane detection on our proposed local and BDD100k datasets in terms of F1 accuracy, IoU, confidence, and FPS. IoU is first computed between predictions and ground truth, and lanes whose IoU exceeds a threshold (0.5) are considered true positives (TP). It is evident in Table 5 that despite CondLaneNet [36] and CLRNet’s [65] high FPS, their accuracy is significantly inferior to PINet’s. Therefore, this gives us the insight that such approaches are not generalizable across different environments. Specifically, they failed when no lines were present in the scene or when other environmental objects obscured the lines. PINet, however, has shown promising results when tested in a new setting, despite some errors that often occur under poor lighting conditions, which we have discussed in 3.1 on how to improve it. Overall, we have selected PINet as our lane detector since it delivers the requisite FPS for the system’s real-time performance while attaining high accuracy in recognizing lanes.

Ablation on Object Detection. Table 6 exhibits the performance of current SOTA approaches for the object detection task on our proposed local dataset and BDD100K. Results indicate that despite the slight difference between approaches, YOLOv5 [27] outperforms all the latest object detection methods in terms of mAP, mAR, and FPS, making it the best choice for our object detection module.

6 Discussion and Limitations

Our experimental results demonstrate the power of our proposed modular system in perceiving the environment and safe driving. The selected networks have been chosen according to the trade-off of accuracy and FPS compared to their counterparts and show high accuracy results. The package FPS can even go higher as the PINet network alone has moderate FPS, and when this network is added to the package, even while the total package stays in real-time, its speed is reduced. Also, we may encounter missing frames on rare occasions, but because it is only one frame in 30 frames, for instance, it is hardly visible. However, by resolving this minor missing frame, the package’s accuracy will improve. In addition, with the development of the package, we can also connect it to the brake and gas pedals. In such a way that after detecting the object at a distance less than a threshold in front of the car, the brake is activated, and the gas pedal is pressed according to the distance from the front cars and the speed limit. It brings us closer to safe driving without the high costs of self-driving vehicles. Overall, the perfect integration of multiple networks and applying appropriate changes to improve their performance, along with the designed GUI, has created a system that gives the necessary warnings to the driver while driving and reduces the risks of driving.

7 Conclusions

In this paper, we have proposed an intelligent modular vision-based system to assist drivers in safe driving by alerting them in critical moments. There are four main stages in our proposed system: lane detection, object and sign detection, sidewalk segmentation, and distance measurement. PINet is used for lane detection and fine-tuned with a combination of our presented local dataset and BDD100K to alleviate its line-free road estimation issue. To prevent PINet from being impacted by turbulence related to the environment, a novel dynamic ROI is applied to the PINet output. Yolov5 has also been utilized to detect 3 class labels and 15 distinct traffic signs. Moreover, having leveraged SGDepth outputs, monocular depth estimation, and segmentation, we have developed a novel hierarchical method to precisely calculate the distance from neighboring vehicles. In addition, a graphical user interface (GUI) is designed to inform the driver about the traffic light’s status, the nearest distance from the adjacent vehicles, and information about traffic signs ahead. Extensive experiments on our local and BDD100K datasets demonstrate that our proposed system performs noticeably well in different environments with different scenarios and conditions and can be an excellent tool to reduce human errors. In the future, we will concentrate on making our system more efficient so that it can operate on embedded boards.

Table 6: Evaluation of different object detection approaches. NVIDIA Tesla T4 is used for evaluation.

		mAP	mAR	FPS
YOLOX [14]	Local	95.90%	89.27%	59.6
	BDD100K	90.72%	96.02%	
YOLOv3 [49]	Local	90.76%	87.94%	37.3
	BDD100K	85.84%	91.43%	
YOLOv4 [7]	Local	93.33%	87.50%	32.1
	BDD100K	88.88%	94.12%	
YOLOv5 [27]	Local	96.01%	90.91%	90.9
	BDD100K	92.42%	96.28%	

References

- [1] Hala Abualsaud, Sean Liu, David B Lu, Kenny Situ, Akshay Rangesh, and Mohan M Trivedi. Laneaf: Robust multi-lane detection with affinity fields. *IEEE Robotics and Automation Letters*, 6(4):7477–7484, 2021.
- [2] Amirhossein Ahmadi, Mojtaba Nabipour, Saman Taheri, Behnam Mohammadi-Ivatloo, and Vahid Vahidinassab. A new false data injection attack detection model for cyberattack resilient energy forecasting. *IEEE Transactions on Industrial Informatics*, 2022.
- [3] Kazi Iftekhar Ahmed. *Modeling drivers' acceleration and lane changing behavior*. PhD thesis, Massachusetts Institute of Technology, 1999.
- [4] C Bradford Barber, David P Dobkin, and Hannu Huuhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software (TOMS)*, 22(4):469–483, 1996.
- [5] Behnam Asghari Beirami and Mehdi Mokhtarzade. Svm classification of hyperspectral images using the combination of spectral bands and moran's i features. In *2017 10th Iranian Conference on Machine Vision and Image Processing (MVIP)*, pages 139–144. IEEE, 2017.
- [6] Abdulbari Bener. The neglected epidemic: road traffic accidents in a developing country, state of qatar. *International journal of injury control and safety promotion*, 12(1):45–47, 2005.
- [7] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. Yolov4: Optimal speed and accuracy of object detection. *arXiv preprint arXiv:2004.10934*, 2020.
- [8] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018.
- [9] Po-Yi Chen, Alexander H Liu, Yen-Cheng Liu, and Yu-Chiang Frank Wang. Towards scene understanding: Unsupervised monocular depth estimation with semantic-aware representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2624–2632, 2019.
- [10] Yuru Chen, Haitao Zhao, Zhengwei Hu, and Jingchao Peng. Attention-based context aggregation network for monocular depth estimation. *International Journal of Machine Learning and Cybernetics*, 12(6):1583–1596, 2021.
- [11] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proc. of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [12] Arturo De La Escalera, Luis E Moreno, Miguel Angel Salichs, and José María Armingol. Road traffic sign detection and classification. *IEEE transactions on industrial electronics*, 44(6):848–859, 1997.
- [13] Kirsten Duivenvoorden. The relationship between traffic volume and road safety on the secondary road network. *SWOV Institute for Road Safety and Research*, 2010.
- [14] Zheng Ge, Songtao Liu, Feng Wang, Zeming Li, and Jian Sun. Yolox: Exceeding yolo series in 2021. *arXiv preprint arXiv:2107.08430*, 2021.
- [15] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015.
- [16] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 580–587, 2014.
- [17] Clément Godard, Oisín Mac Aodha, and Gabriel J. Brostow. Unsupervised Monocular Depth Estimation With Left-Right Consistency. In *Proc. of CVPR*, pages 270–279, Honolulu, HI, USA, July 2017.
- [18] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging Into Self-Supervised Monocular Depth Estimation. In *Proc. of ICCV*, pages 3828–3838, Seoul, Korea, October 2019.
- [19] Robert M. Haralick, Stanley R. Sternberg, and Xinhua Zhuang. Image analysis using mathematical morphology. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-9(4):532–550, 1987.
- [20] Malik Haris and Adam Glowacz. Lane line detection based on object feature distillation. *Electronics*, 10(9):1102, 2021.

- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- [22] Paniz Hosseini. Deep-learning neural network prediction of a solar-based absorption chiller cooling system performance using waste heat. *Sustainable Energy Technologies and Assessments*, 53:102683, 2022.
- [23] Paniz Hosseini, Sadegh Nikbakht Naserabad, Amir H Keshavarzzadeh, and Najmeh Ansari. Artificial intelligence-based tri-objective optimization of different demand load patterns on the optimal sizing of a smart educational buildings. *International Journal of Energy Research*, 2022.
- [24] Shantanu Ingle and Madhuri Phute. Tesla autopilot: semi autonomous driving, an uptick for future autonomy. *International Research Journal of Engineering and Technology*, 3(9):369–372, 2016.
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [26] Licheng Jiao, Fan Zhang, Fang Liu, Shuyuan Yang, Lingling Li, Zhixi Feng, and Rong Qu. A survey of deep learning-based object detection. *IEEE access*, 7:128837–128868, 2019.
- [27] Glenn Jocher, Ayush Chaurasia, Alex Stoken, Jirka Borovec, NanoCode012, Yonghye Kwon, TaoXie, Jiacong Fang, imyhxy, Kalen Michael, Lorna, Abhiram V, Diego Montes, Jebastin Nadar, Laughing, tkianai, yxNONG, Piotr Skalski, Zhiqiang Wang, Adam Hogan, Cristi Fati, Lorenzo Mammanna, AlexWang1900, Deep Patel, Ding Yiwei, Felix You, Jan Hajek, Laurentiu Diaconu, and Mai Thanh Minh. ultralytics/yolov5: v6.1 - TensorRT, TensorFlow Edge TPU and OpenVINO Export and Inference, February 2022.
- [28] Jong Bae Kim. Efficient vehicle detection and distance estimation based on aggregated channel features and inverse perspective mapping from a single camera. *Symmetry*, 11(10):1205, 2019.
- [29] Stuart T Klapp. Motor response programming during simple choice reaction time: The role of practice. *Journal of Experimental Psychology: Human perception and performance*, 21(5):1015, 1995.
- [30] Marvin Klingner, Jan-Aike Termöhlen, Jonas Mikolajczyk, and Tim Fingscheidt. Self-supervised monocular depth estimation: Solving the dynamic object problem by semantic guidance. In *European Conference on Computer Vision*, pages 582–600. Springer, 2020.
- [31] Yeongmin Ko, Younkwan Lee, Shoaib Azam, Farzeen Munir, Moongu Jeon, and Witold Pedrycz. Key points estimation and point instance segmentation approach for lane detection. *IEEE Transactions on Intelligent Transportation Systems*, 2021.
- [32] Hui Kong, Jean-Yves Audibert, and Jean Ponce. General road detection from a single image. *IEEE Transactions on Image Processing*, 19(8):2211–2220, 2010.
- [33] Minhyeok Lee, Sangwon Hwang, Chaewon Park, and Sangyoum Lee. Edgeconv with attention module for monocular depth estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2858–2867, 2022.
- [34] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [35] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [36] Lizhe Liu, Xiaohao Chen, Siyu Zhu, and Ping Tan. Condlanenet: a top-to-down lane detection framework based on conditional convolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3773–3782, 2021.
- [37] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016.
- [38] Weifeng Liu, Zhenqing Zhang, Shuying Li, and Dapeng Tao. Road detection by using a generalized hough transform. *Remote Sensing*, 9(6):590, 2017.
- [39] S Mahdi H Miangoleh, Sebastian Dille, Long Mai, Sylvain Paris, and Yagiz Aksoy. Boosting monocular depth estimation models to high-resolution via content-adaptive multi-resolution merging. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9685–9694, 2021.

- [40] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European conference on computer vision*, pages 483–499. Springer, 2016.
- [41] Abbas Omid, Amirhossein Heydarian, Aida Mohammadshahi, Behnam Asghari Beirami, and Farzan Haddadi. An embedded deep learning-based package for traffic law enforcement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 262–271, 2021.
- [42] Ekaterina Panfilova, Oleg S Shipitko, and Irina Kunina. Fast hough transform-based road markings detection for autonomous vehicle. In *Thirteenth International Conference on Machine Vision*, volume 11605, page 116052B. International Society for Optics and Photonics, 2021.
- [43] Adam Paszke, Abhishek Chaurasia, Sangpil Kim, and Eugenio Culurciello. ENet: A Deep Neural Network Architecture for Real-Time Semantic Segmentation. *arXiv*, June 2016. (arXiv:1606.02147).
- [44] Zhan Qu, Huan Jin, Yang Zhou, Zhen Yang, and Wei Zhang. Focus on local: Detecting lane marker from bottom up via key point. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14122–14130, 2021.
- [45] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10428–10436, 2020.
- [46] Prakash Ranjitkar, Takashi Nakatsuji, Yoichi Azuta, and Gemunu Senadeera Gurusinghe. Stability analysis based on instantaneous driving behavior using car-following data. *Transportation Research Record*, 1852(1):140–151, 2003.
- [47] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [48] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016.
- [49] Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- [50] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015.
- [51] Obeid Sharifi, Mehdi Mokhtarzade, and Behnam Asghari Beirami. A deep convolutional neural network based on local binary patterns of gabor features for classification of hyperspectral images. In *2020 International Conference on Machine Vision and Image Processing (MVIP)*, pages 1–5. IEEE, 2020.
- [52] Suvarna Shirke and R Udayakumar. Lane datasets for lane detection. In *2019 International Conference on Communication and Signal Processing (ICCSPP)*, pages 0792–0796. IEEE, 2019.
- [53] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7262–7272, 2021.
- [54] Samia Sultana and Boshir Ahmed. Robust nighttime road lane line detection using bilateral filter and sagc under challenging conditions. In *2021 IEEE 13th International Conference on Computer Research and Development (ICCRD)*, pages 137–143. IEEE, 2021.
- [55] Lucas Tabelini, Rodrigo Berriel, Thiago M Paixao, Claudine Badue, Alberto F De Souza, and Thiago Oliveira-Santos. Keep your eyes on the lane: Real-time attention-guided lane detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 294–302, 2021.
- [56] Domen Tabernik and Danijel Skočaj. Deep Learning for Large-Scale Traffic-Sign Detection and Recognition. *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- [57] S. Tuohy, D. O’Cualain, E. Jones, and M. Glavin. Distance determination for an automobile environment using inverse perspective mapping in opencv. In *IET Irish Signals and Systems Conference (ISSC 2010)*, pages 100–105, 2010.
- [58] Lijun Wang, Jianming Zhang, Oliver Wang, Zhe Lin, and Huchuan Lu. Sdc-depth: Semantic divide-and-conquer network for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 541–550, 2020.

- [59] Ying Sheng Wang, Yuhan Qi, and Yi Man. An improved hough transform method for detecting forward vehicle and lane in road. *Journal of Physics: Conference Series*, 1757(1):012082, jan 2021.
- [60] Dong Wu, Manwen Liao, Weitian Zhang, and Xinggong Wang. Yolop: You only look once for panoptic driving perception. *arXiv preprint arXiv:2108.11250*, 2021.
- [61] Sitong Wu, Tianyi Wu, Fangjian Lin, Shengwei Tian, and Guodong Guo. Fully transformer networks for semantic image segmentation. *arXiv preprint arXiv:2106.04108*, 2021.
- [62] Guorun Yang, Hengshuang Zhao, Jianping Shi, Zhidong Deng, and Jiaya Jia. Segstereo: Exploiting semantic information for disparity estimation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 636–651, 2018.
- [63] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. Bdd100k: A diverse driving dataset for heterogeneous multitask learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2636–2645, 2020.
- [64] Zhenyu Zhang, Zhen Cui, Chunyan Xu, Zequn Jie, Xiang Li, and Jian Yang. Joint task-recursive learning for semantic segmentation and depth estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 235–251, 2018.
- [65] Tu Zheng, Yifei Huang, Yang Liu, Wenjian Tang, Zheng Yang, Deng Cai, and Xiaofei He. Clnet: Cross layer refinement network for lane detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 898–907, 2022.
- [66] Zhe Zhu, Dun Liang, Songhai Zhang, Xiaolei Huang, Baoli Li, and Shimin Hu. Traffic-sign detection and classification in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.