A Review of Non-Lane Road Marking Detection and Recognition

Adam Morrissett, Sherif Abdelwahed
Department of Electrical and Computer Engineering
Virginia Commonwealth University
Richmond, Virginia 23284–3068
Emails: morrissettal2@vcu.edu, sabdelwahed@vcu.edu

***Abstract*—Environment perception is a critical function used by driving automation systems, or self-driving cars, for detecting objects such as obstacles, lane markings, and road signs. In order to replace human drivers, self-driving cars will need to safely operate in parking lots, private roads, underground, or any other environment with poor GPS signals or uncharted infrastructure. While much attention has been spent on recognizing lane markings, non-lane road markings have received considerably less attention. Current perception systems can recognize only a small subset of markings and often only under favorable weather conditions. This limitation is exacerbated by the current quality of scene segmentation data sets. Only a select few of existing data sets have annotations for non-lane road markings, and the ones that do only have them for a small number of marking types. Additionally most of the data sets were generated under one type of driving condition. Finally, it is difficult to determine if current recognition systems can satisfy real-time requirements. This paper investigates the current limitations and challenges for non-lane road marking detection and recognition including recognition capabilities, data set quality, and inference times.**

I. Introduction

Signs and road markings visually convey the laws and regulations governing road networks; they can also indicate upcoming hazards or points of interest. For example, raised signs can display speed limits, traffic flow, and obstructed or "blind" turns [1], [2]. Additionally, painted solid and striped lines indicate lane boundaries. Other road markings, sometimes termed *non-lane* [3], [4] or *symbolic* [5]–[7] road markings, convey similar information to raised signs through text and/or symbols. This can include speed limits, legal turning maneuvers, and lane restrictions [2], [8]. These types of markings are categorized separately from lane road markings because their immediate purpose is not to assist drivers with staying within road lane bounds.

Obeying traffic regulations and heeding warning signs promotes a safe and predictable driving environment. Therefore, driving automation systems [9] must conform to the same conventions as human drivers if they will be operating in shared environments. This is especially critical on urban roads, which are often shared with cyclists and pedestrians. Besides non-lane road markings, driving automation systems can receive road information from raised signs [1] and map metadata [10]. While raised signs are often abundant, and maps are increasingly more detailed, this information is unavailable in certain areas. Private parking lots and roads are often uncharted, and cellular networks and global positioning systems (GPSs) are unreliable in concrete structures, such as parking garages. Additionally, some urban areas use only road markings to convey road information [11]. In these cases, driving automation systems are forced to rely on non-lane road markings to indicate appropriate driving maneuvers.

Human drivers can easily recognize and interpret a variety of road signs and markings, but this task can be challenging for driving automation systems. Obstructed or degraded markings and inconsistent light conditions can hinder detection and recognition accuracy. Non-lane road markings also vary from location to location, making it more difficult for detection and recognition systems to generalize. Furthermore, detection and recognition systems need to operate quickly and predictably in order to satisfy real-time constraints.

Despite its necessity, detecting and recognizing non-lane road markings remains a relatively open problem compared to lane-marking detection. Current machine-learning-based approaches can achieve high recognition accuracy but only for a small subset of marking classes. Many of these approaches have been tested only under a small number of different driving conditions; some have been tested under only one. Furthermore, many existing works do not disclose the runtime performances of their systems, making it difficult to determine the current state of the art in that aspect.

Data set quality is another factor limiting detection and recognition systems. Out the 11 data sets mentioned in this paper, only 4 include annotations for non-lane road markings [7], [12]–[14]; the others do not distinguish between the road surface and non-lane road markings. These 4, however, annotate only a small number of markings, much less than the number of existing ones. Also, most of the data sets were generated under only sunny, favorable weather conditions.

To facilitate research related to non-lane road marking detection and recognition, this paper provides a survey on the current state of the art including detection and recognition capabilities, data set quality, and inference times. The rest of the paper is structured as follows. Section II formulates the problem and the typical pipelines used in detection and recognition systems. Section III discusses conventional approaches toward non-lane road marking detection and recognition.

(a) (b) (c)

Fig. 1. Non-lane road marking examples. (a) Railroad crossing [12]. (b) "Keep clear" text [12]. (c) Front arrow and "bridge" text [13].

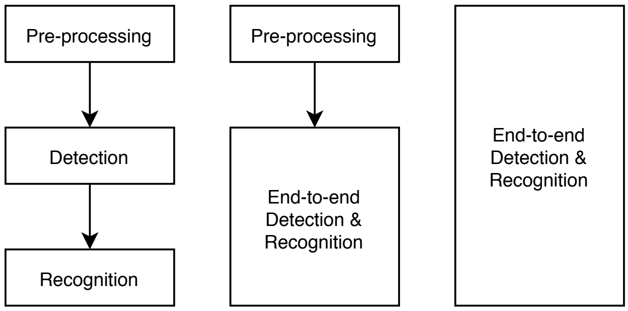

Fig. 2. Detection and recognition pipelines. (left to right) Conventional pipeline, partial end-to-end pipeline, full end-to-end pipeline.

Section IV discusses end-to-end detection and recognition systems. Section V discusses the typical data sets used in non-lane road marking detection and recognition. Section VI presents open problems and suggests future research directions. Finally, Section VII concludes the paper.

II. Problem Formulation

Road networks throughout the world are supported by a variety of information encoded in various formats; one of these formats is non-lane road markings. Specifically, non-lane road markings are markings containing text and/or symbols that are painted on roads to convey information to drivers. This information can include, but is not limited to, points of interest, speed limits, or available maneuvering options.

Non-lane road marking types and specifications vary by location. For example, Virginia actively uses at least 25 different symbolic markings including turning arrows, special lane indicators, yield symbols, and railroad crossings among other symbols [15]. Note this list does not include text-based markings or a combination of text and symbolic markings. Fig. 1 shows examples of different types of markings taken from the Road Marking [12] and Tsinghua Road Marking [13] data sets, explained later.

Environment perception is a challenging task for humans. The driving environment is cluttered with distractions such as advertising billboards, varying scenery, and interesting buildings. While they may be pleasing to look at, these features can distract drivers from numerous important signs that inform them about the road. The driving scene is often further cluttered in urban environments where visual distractions are significantly denser. To make driving even more challenging, drivers need to correctly interpret all of these signs continuously; new signs are always approaching with new information. Drivers are also expected to accurately and quickly read signs and markings in adverse conditions, such as fog, heavy rain, darkness, or any combination.

Perception is even more complicated for computers because in addition to challenges in capturing images, they do not possess intuition. Images can be noisy, especially in low-light conditions. They can also be blurry if the camera is not capturing fast enough (i.e., the shutter speed is too slow) with respect to movement speed. Illumination changes affecting human perception are exacerbated in camera systems; what may be difficult for a human driver to see may be impossible for a camera to see. Marking meaning or applicability can also change depending on context. Despite these numerous challenges, recent works in computer vision and machine learning have shown promising results.

Many works using traditional computer vision frameworks followed a similar pipeline consisting of pre-processing, marking detection, and marking recognition. As machine learning techniques have become popular, some researchers substituted portions of this pipeline for machine learning algorithms, creating a partial end-to-end system. Other researchers have gone further and replaced the whole pipeline with machine learning algorithms creating a full end-to-end system. Fig. 2 provides a visual representation of these pipelines. The rest of this paper will discuss the different approaches in relation to this pipeline.

III. Conventional Detection and Recognition

1) Pre-Processing: We define pre-processing as image transformations that do not explicitly function as detection or recognition sub-systems, but rather remove noise or other distractions to aid in these processes. Some of the types of pre-processing commonly employed are region of interest (ROI) reduction, thresholding, and inverse perspective mapping (IPM).

One simple method of reducing the size of ROIs is by removing areas of the scene in which there are no road markings, such as above the horizon [3], [16], [17]. By assuming all markings will be in the lower half of the image, the ROI size can be easily reduced without any complex operations. Two problems with this is that it requires precise positioning of the camera to avoid losing relevant data and that it still leaves unnecessarily large ROIs.

A more sophisticated technique for eliminating areas above the horizon is implementing an automatic horizon or vanishing point detection mechanism [18], [19]. Other researchers limit the ROIs to the current lane of the ego-vehicle [17], [20]. However, this also restricts marking detection to the current lane even though other information, such as legal turning maneuvers, could be displayed in adjacent lanes. Some works limit the ROI to all detected lanes to allow for a better understanding of the entire roadway [21], [22]. However, all of these methods require vehicles to have lane detection functionality and for this functionality to be operational.

Another pre-processing technique called thresholding is used to remove any pixels that are not part of the marking candidates. One thresholding method, called the Otsu method, is used to filter out pixels that are not between lane lines [22]. Median thresholding, which uses the median value of the entire image row, is another type [7]. A modified version of median local thresholding, called 43rd percentile thresholding, allows markings to be filtered based on size [23]. Other types of thresholding include high brightness filtering [6], top-hat filtering [24], and adaptive image thresholding [3], [19].

Finally, inverse perspective mapping (IPM) has been an incredibly popular pre-processing step as it is used by almost all works implementing traditional or partially end-to-end pipelines. This technique allows perspective distortions from the camera to be removed and for a top-down version of the image to be generated, thus allowing for easier and more accurate road marking detection and recognition. However, one problem associated with this technique is that it requires knowledge about the camera, and it requires camera calibration [25].

2) Detection: After the image has been pre-processed, non-lane road marking candidates are detected and separated from the background. A variety of techniques have been used for marking candidate selection. Back projection can be used to determine the overall size of each marking, and is a popular choice for template-based classifiers [6], [21], [22].

Connected component analysis and principal component analysis (PCA), have been widely used techniques for detecting and extracting potential marking candidates. Connected component analysis has been used by [20], [21], [23], and PCA has been used by [3], [26]. A variant of connected component analysis, called marker-based connected component analysis has also been used [24], [27]. Additionally, connected component analysis has been coupled with a support vector machine (SVM) to further refine candidate selection based on histogram of gradient (HOG) features [28]. Maximally stable extremal regions (MSERs) have also been proposed to distinguish between marking candidates and the rest of the roadway [12], [29], [30].

Machine learning has also been used for marking candidate selection. In [31], the binarized normed gradient (BING) method is used in conjunction with an SVM to detect possible road markings. Researchers in [18] used an SVM to improve their marker-based watershed algorithm. A deconvolution network is used in [7] to detect and classify pixels in the frame according a particular category, such as road, lane, or sky.

Binarization has been a common method of transforming any filtered images into a binary image before being sent to the recognition module [6], [18], [21], [28]. For many machine-learning-based detection and recognition systems, one of the final stages of detection has been to generate histograms of gradients (HOGs) [6], [12], [24], [28], [29]. HOGs are a way to represent the features of an object by deriving the gradient of groups of pixels in all directions.

3) Recognition: After marking candidates have been detected, they can be classified. A variety of methods have been proposed to do so, with the most popular being template matching, SVMs, and variations of neural networks.

Originally, marking detection and recognition was performed using template matching [12], [21], [23], [25], [26]. Template, or geometric pattern, matching involves creating a representation of each marking class to create a database which is then used to compare unknown marking candidates. These templates have been created based on national standards for markings, and the classifier determines the validity of a candidate based on how well it matches the database templates [32]. Variations of this method include prototype fitting with arc splines [20] and pattern matching of geometric sub-components of marking candidates [17].

The most significant problem with the pattern matching approaches is the dependence on national standards. Because the templates are generated from the specifications listed in roadway standards documents, they are effective only in locations that implement the assumed standards. Additionally, the deterioration, obstruction, or design changes of these markings make pattern matching approaches difficult to generalize. These systems also lack the accuracy required for reliable use with many achieving around 90% or less accuracy.

Eventually, SVMs and other machine learning techniques started becoming more popular. A multi-class SVM classifier in which a single SVM was dedicated to detecting and classifying a single marking class was proposed in [22], [28]. Similarly, a decision tree was proposed by [33]. In [31], a PCANet-based classifier was proposed in which a PCANet [34] was used to filter out false marking candidates, and an SVM was used to classify the remaining true candidates. A cascaded classifier using AdaBoost was used in [24], where Haar-like descriptors were used for coarse detection and an extreme learning machine (ELM) was used for fine recognition. However, this system was only able to classify arrow and diamond markings. Another cascaded classifier was proposed in [6] in which a two-class total error rate (TER)-based classifier was used to reject non-marking candidates, and a multi-class TER-based classifier was used for marking candidate recognition.

Some researchers began to add functionality to also detect and classify text-based markings. Researchers in [29] used a multi-level classifier with an optical character recognition (OCR) classifier for text-based markings and an SVM for symbol-based markings. In turn, those in [18] used an OCR-based classifier for both text- and symbols-based markings with the Tesseract library.

IV. End-to-End Detection and Recognition

More recent approaches have used variants of neural networks and other deep learning methods, specifically convolutional neural networks (CNNs). In [19], contour representations of marking candidates were passed through an artificial neural network (ANN). This network contained a single hidden layer and used back-propagation for training.

TABLE I
CLASSIFICATION ABILITIES SUMMARY

Author	Year	FA	LA	RA	LRA	FLA	FRA	FLRA	*STOP*	*BIKE*	*PED*	*XING*	*35*	*40*	RR	Bike	Others	Classes	F_1 Score	Average Accuracy
Vokhidov *et al.* [35]	2016	✓	✓	✓		✓	✓	✓										6	0.9999	0.9999
Ahmad *et al.* [5]	2017	✓	✓	✓					✓	✓	✓	✓	✓	✓	✓			10		0.9905
Bailo *et al.* [30]	2017	✓	✓	✓					✓		✓	✓	✓	✓	✓		✓	10		0.9080
Lim *et al.* (Sym-1) [7]	2018																✓	14	0.9080	
Lim *et al.* (Sym-2) [7]	2018																✓	14	0.9365	
Hoang *et al.* [36]	2019	✓	✓	✓	✓	✓	✓	✓								✓		8	0.9610	0.9280
Lee *et al.* [37]	2019	✓	✓	✓					✓		✓	✓	✓	✓	✓		✓	10		0.9530

It was also able to detect and classify both symbols and text. In [5], a CNN based on LeNet [38] was used to classify both symbol- and text-based road markings. The CNN was trained using the Road Marking data set [12], and different architecture variations were compared. Researchers in [30] proposed and compared the performance of three different classifiers: PCANet and an SVM, PCANet and logistic regression, and a shallow CNN. It was also trained using the Road Marking data set [12]. A cascade classifier consisting of a CNN and an SVM was proposed by [7]. The CNN was used for marking detection and recognition, and the SVM was used for lane detection. The authors evaluated the classifier using their own data set [7]. A CNN was also used in [35], which was trained using multiple data sets [12], [39]–[42]. Unlike other learning-based systems, this one was developed with an emphasis on detecting and recognizing damaged markings. Additionally, a fully convolutional neural network (FCNN) was proposed by [16] for detecting and segmenting road markings from the rest of the driving scene. It should be noted that this work only proposed an marking detector for use in a larger system; it was not used for recognition.

Common limitations for CNN-based detection and recognition systems include low classification accuracy when presented with markings that are blurry, deteriorated, poorly-illuminated, or otherwise distorted [30], [35]–[37]. To address this limitation, Lee *et al.* [37] used a two-stage system structured similarly to a generative adversarial network (GAN). The first stage of their network, called the *generator network*, used a fully-convolutional auto-encoder network to de-blur images while preserving salient features. The second stage, called the *discriminator network*, used convolutional layers to determine the marking class and reject false positives. Like many others, the system was evaluated using the Road Marking data set.

Table I summarizes the classification abilities, F_1 scores (if available), and overall accuracies (if available) for different methods. The abbreviations are defined as follows: front arrow (FA), left arrow (LA), right arrow (RA), left-right arrow (LRA), front-left arrow (FLA), front-right arrow (FRA), front-left-right arrow (FLRA), *STOP* text, *BIKE* text, *PED* (short for *pedestrian*) text, *XING* (short for *crossing*) text, *35* (a speed limit) text, *45* (another speed limit) text, rail road (RR) crossing symbol, and a bicycle symbol. Note that in the entry for Lim *et al.* [7], the specific classes are not populated because the authors did not provide a list of specific classes; they only provided the total count.

While Table I seems to indicate that non-lane road marking detection and recognition is a solved problem, it should be emphasized that these results were achieved within a specific environment. The presented systems are able to classify only a small subset of commonly-used markings. None of the approaches mentioned discuss evaluation results for marking detection and recognition in adverse weather conditions, such as rain. Additionally, some methods performed poorly when presented with damaged or malformed markings [30], [36].

V. DATA SETS

As mentioned in a recent survey from Kang *et al.* [43], a variety of public data sets exist to facilitate driving automation system research. Two highly popular ones are the CityScapes [44] and Karlsruhe Institute of Technology and Toyota Technological Institute (KITTI) [39], [40] data sets. The CityScapes data set was created to assist with semantic segmentation and driving scene understanding [44]. Images were generated using a stereo camera, and the driving scenes were annotated according to 30 different classes (e.g., *road*, *sidewalk*, *car*). Additionally, the classes were collected into 8 different categories, such as *flat*, *vehicle*, and *sky*. The data set also includes vehicle odometry, outside temperature, and GPS data.

The KITTI data set was designed to facilitate general computer vision and autonomous vehicle research [39]. The data acquisition vehicle was equipped with both grey-scale and color stereo cameras, a 3D lidar, and a GPS and IMU navigation system. Images were taken in city, residential, campus, and other driving environments. The images also include annotations for different objects, such as cars, trucks, and people. The KITT-ROAD data set [40], includes road area and ego lane annotations for 600 images taken from the original KITTI data set. It includes images taken from urban environments with no lane markings, two-way lane markings, and multi-lane markings.

The Oxford RobotCar [45], Berkeley Deep Drive 100K (BDD100K) [46], and Málaga [41], [42] data sets are also popular. The Oxford RobotCar data set focused on providing data for the same place over a period of time. It is unique compared to the other data sets because it's focus is on providing data of the same route over long periods of time.

By continually traveling along the same route, the authors were able to collect data for a variety of weather conditions and environment changes.

The BDD100K data set provides data for object detection, lane marking recognition, drivable area detection, and semantic segmentation. It is composed of crowd-sourced videos taken under 6 different weather conditions (including snow and rain), 3 different times of day, and in urban, residential, and highway environments. The data set also contains GPS and IMU data. Annotations include bounding boxes for 10 different object categories, labels for 22 different semantic segmentation classes, and labels for 8 different lane markings. It should be noted that the data set does not include labels for non-lane road markings.

The Málaga data sets were developed for simultaneous localization and mapping (SLAM) research. The data acquisition vehicles for both data sets were equipped with cameras, 2D lidars, a GPS, and an IMU. The Málaga 2009 data set [41] contains images of the driving scene of parking lot and college campus environments, but none of them include annotations. Málaga Urban [42] data set contains images of the driving scene in parking lots, urban, and suburban driving environments, but it also does not include annotations.

While all of previously-mentioned data sets have image frames capturing non-lane road markings, they do not provide annotations for them. Therefore, researchers have collected other data sets purpose-built for non-lane road marking detection and recognition. It should be mentioned that the researchers in [35] used the KITTI and Málaga data sets, but their non-lane road marking annotations have not been publicized to the best of our knowledge.

One of these purpose-built data sets is the Road Marking data set [12]. This data set contains 1 443 images captured from a front-facing camera mounted on the roof of a vehicle, and the images were taken in a variety of conditions. Images were annotated by hand with ground-truth labels for a variety of arrow configurations and a variety of text-based markings.

More recently, researchers in [13] presented the Tsinghua Road Marking (TRoM) data set. This data set contains 712 images and includes annotations for 6 different marking classes. While it does not have the highest number of class types, it does have 6 different weather conditions, the most out of all the non-lane marking data sets and the second most out of all the data sets mentioned in this paper.

Most recently, researchers in [14] presented the ApolloScape data set. The researchers used a high-resolution camera to capture 165 949 images, which were annotated to include 20 different marking classes.

Another recent data set comes from Lim *et al.* [7]. Their data set is composed of two training data sets, T-Sym and T-Lane, and four testing data sets, Lane-1, Lane-2, Sym-1, and Sym-2. The data sets with *Sym* in their name are used for non-lane, or symbolic, road marking recognition, and the ones with *Lane* in their name are used for lane marking recognition. The combined symbolic marking data set contains 36 000 images with 14 different marking classes. It should be noted that the authors do not mention what the marking classes are.

Table V provides a comparison of the different data sets and helps emphasize their limitations as a whole. The last column indicates the number of weather conditions contained within the data set (sunny, overcast, rainy, nighttime, etc.). As shown by the table, only 4 of the surveyed data sets contain annotations for non-lane road markings. Despite containing images of them, the other data sets do not have annotations for these markings. Furthermore, many of the data sets include only one type of weather condition.

VI. Open Problems and Future Directions

Learning-based marking detection and recognition methods show improved accuracy and robustness compared to conventional methods, but, like conventional methods, they too can only classify a limited number of markings. As mentioned previously, none of the reviewed learning-based systems show evaluation results for adverse weather conditions, such as rain, snow, or fog. Some systems also demonstrated poor performance when classifying deteriorated or malformed markings. Human drivers commonly encounter these conditions, so detection and recognition systems will need to perform well under these conditions too.

In addition to limited classification abilities, many learning-based methods are too slow to be implemented in real-time driving automation systems. To the best of our knowledge, only Hoang *et al.* [36] has implemented a non-lane road marking detection and recognition system with an emphasis on real-time performance. Similar systems have been recently proposed for general semantic segmentation, but they have not yet been implemented specifically for non-lane road marking detection and recognition [47], [48]. This could be a promising research direction.

While current machine-learning-based recognition systems are incredibly accurate, especially compared to other methods, significant work still remains before they are ready for production in autonomous vehicles. The most challenging of which is the current state of data sets. As mentioned by [49], [50], more data sets are needed to improve the overall performance of marking detection and recognition systems and to prevent over-fitting of training data. A larger variety of data sets are also needed to provide training data that encompasses a wide variety of marking types from different locations and in different environmental conditions. Additionally, data sets need to include images for a variety of weather conditions, such as rain, storms, and snow; drivers commonly experience all of these conditions. Improving existing data sets would be a significant contribution to the research field.

Furthermore, data sets will need to expand to include more categories of non-lane road markings. The current amount of road marking classes that can be recognized is significantly less than the number used of roads. The Virginia Department of Transportation (VDOT) lists over 25 symbolic-only road markings and over 9 alpha-numeric-based road markings [51]. Most current data sets include less than half of these.

Time and resource requirements for these systems is another problem. Many of the end-to-end implementations

TABLE II
DATA SET SUMMARY

Data Set	Year	Images	# Markings	# Weather Conditions
Road Marking [12]	2012	1 443	10	4
TRoM [13]	2017	712	6	6
ApolloScape [14]	2018	165 949	20	1
T-Sym, Sym-1, and Sym-2 [7]	2018	36 200	14	1
Oxford RobotCar [45]	2017	20 000 000[1]	0	7
BDD100K [46]	2018	100 000[2]	0	9
KITTI [39]	2013	44 054	0	1
KITTI-ROAD [40]	2013	600	0	1
Málaga 2009 [41]	2009	9 998	0	1
Málaga Urban [42]	2009	113 082	0	1
Cityscapes [44]	2016	25 000	0	1

[1] Image count is an approximation based on the data set website (https://robotcar-dataset.robots.ox.ac.uk/documentation/)
[2] Image count is based on the number of lane marking annotations

mentioned previously did not disclose the hardware used for testing. Those that did describe the testing hardware used high-performance computers and graphics processing units (GPUs) such as the NVIDIA GeForce GTX Titan X [35], [37] and NVIDIA GeForce GTX 1070 [36]. Furthermore, the processing time for each frame with some implementations is incredibly slow at only a few frames per second [16]. This may not be feasible if implemented in vehicles because the system would not be able to successfully operate, especially with the vehicle traveling at speed. If implemented in a vehicle, these systems would have time and resource constraints that could reduce their viability. For comparison, the NVIDIA DRIVE AGX Xavier system (a purpose-built system-on-chip for self-driving cars) has only 1024 CUDA cores while the GTX Titan X and GTX 1070 have 3072 and 1920 cores, respectively [52]–[54].

As autonomous vehicles become more sophisticated, non-lane road marking detection and recognition systems will need to operate using shared resources while still maintaining performance goals, which may not be possible in their current state. In the future, there will have to be a trade-off between performance and accuracy. To help emphasize the importance of efficiency of their detection and recognition systems, we encourage researchers to include system utilization metrics as part of their experimentation results.

VII. CONCLUSION

Non-lane road markings convey important traffic information that is sometimes unavailable in other forms. Accurately detecting and recognizing these markings is critical for improving driving automation systems' abilities. Current implementations have limited recognition abilities that are further limited when dealing with damaged or obstructed markings and changing light conditions. Inference times are too large to be used in real-time applications, especially if computing hardware is shared with other perception or control algorithms. Furthermore, data sets contain only a limited number of markings and weather conditions, thus hindering development. In this paper, we discussed the current state of the art and limitations for non-lane road marking detection and recognition systems. Additionally, we discussed current limitations with data sets, and we suggested future research directions based on currently open problems.

REFERENCES

[1] Virginia Department of Transportation, *Virginia Standard Highway Signs*, 2011th ed. Richmond, VA: Virginia Department of Transportation, 2015.

[2] Department for Transport, *Know Your Traffic Signs: Official Edition*, 5th ed. Norwich, UK: The Stationary Office (TSO), 2007.

[3] A. Gupta and A. Choudhary, "A Framework for Camera based Real-Time Lane and Road Surface Marking Detection and Recognition," *IEEE Transactions on Intelligent Vehicles*, vol. 3, no. 4, pp. 476–485, 2018.

[4] D. Hyeon, S. Lee, S. Jung, S. W. Kim, and S. W. Seo, "Robust Road Marking Detection Using Convex Grouping Method in Around-View Monitoring System," in *2016 IEEE Intelligent Vehicles Symposium (IV)*. Gothenburg, Sweden: IEEE, 2016, pp. 1004–1009.

[5] T. Ahmad, D. Ilstrup, E. Emami, and G. Bebis, "Symbolic Road Marking Recognition Using Convolutional Neural Networks," in *2017 IEEE Intelligent Vehicles Symposium (IV)*. Los Angeles, CA, USA: IEEE, 2017, pp. 1428–1433.

[6] J. K. Suhr and H. G. Jung, "Fast Symbolic Road Marking and Stop-line Detection for Vehicle Localization," in *2015 IEEE Intelligent Vehicles Symposium (IV)*. Seoul, South Korea: IEEE, 2015, pp. 186–191.

[7] K. Lim, Y. Hong, M. Ki, Y. Choi, and H. Byun, "Vision-Based Recognition of Road Regulation for Intelligent Vehicle," in *2018 IEEE Intelligent Vehicles Symposium (IV)*. Changshu, China: IEEE, 2018, pp. 1418–1425.

[8] Virginia Department of Transportation, *Virginia Supplement to the 2009 Manual on Uniform Traffic Control Devices for Streets and Highways*, 2011th ed. Richmond, VA: Virginia Department of Transportation, 2013.

[9] SAE International, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," SAE International, Tech. Rep., 2018.

[10] Google Developers, "Google Maps Platform Documentation: Roads API." [Online]. Available: https://developers.google.com/maps/documentation/roads/intro

[11] Federal Highway Administration, *Manual on Uniform Traffic Control Devices for Streets and Highways*, 2009th ed. Federal Highway Administration, 2012.

[12] T. Wu and A. Ranganathan, "A Practical System for Road Marking Detection and Recognition," in *2012 IEEE Intelligent Vehicles Symposium*. Alcala de Henares, Spain: IEEE, 2012, pp. 25–30.

[13] X. Liu, Z. Deng, H. Lu, and L. Cao, "Benchmark for Road Marking Detection: Dataset Specification and Performance Baseline," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. Yokohama, Japan: IEEE, 2017.

[14] X. Huang, X. Cheng, Q. Geng, B. Cao, D. Zhou, P. Wang, Y. Lin, and R. Yang, "The ApolloScape Dataset for Autonomous Driving," in *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Salt Lake City, UT, USA: IEEE, 2018, pp. 1067–1073.

[15] Virginia Department of Transportation, “Pavement Word, Symbol, and Arrow Markings,” Virginia Department of Transportation, Tech. Rep., 2015.

[16] L. R. T. Horita and V. Grassi, “Employing a Fully Convolutional Neural Network for Road Marking Detection,” in *2017 Latin American Robotics Symposium (LARS) and 2017 Brazilian Symposium on Robotics (SBR)*. Curitiba, Brazil: IEEE, 2017.

[17] S. Suchitra, R. K. Satzoda, and T. Srikanthan, “Detection & classification of arrow markings on roads using signed edge signatures,” *IEEE Intelligent Vehicles Symposium, Proceedings*, vol. 1, pp. 796–801, 2012.

[18] M. Schreiber, F. Poggenhans, and C. Stiller, “Detecting symbols on road surface for mapping and localization using OCR,” *2014 17th IEEE International Conference on Intelligent Transportation Systems, ITSC 2014*, pp. 597–602, 2014.

[19] A. Kheyrollahi and T. P. Breckon, “Automatic real-time road marking recognition using a feature driven approach,” *Machine Vision and Applications*, vol. 23, no. 1, pp. 123–133, 2012.

[20] G. Maier, S. Pangerl, and A. Schindler, “Real-time detection and classification of arrow markings using curve-based prototype fitting,” *IEEE Intelligent Vehicles Symposium, Proceedings*, no. Iv, pp. 442–447, 2011.

[21] S. Vacek, C. Schimmel, and R. Dillmann, “Road-marking analysis for autonomous vehicle guidance,” in *3rd European Conference on Mobile Robots (EMCR)*, Freiburg, Germany, 2007.

[22] N. Wang, W. Liu, C. Zhang, H. Yuan, and J. Liu, “The Detection and Recognition of Arrow Markings Recognition Based on Monocular Vision,” in *2009 Chinese Control and Decision Conference*. Guilin, China: IEEE, 2009, pp. 4380–4386.

[23] P. Foucher, Y. Sebsadji, J. P. Tarel, P. Charbonnier, and P. Nicolle, “Detection and Recognition of Urban Road Markings Using Images,” in *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*. Washington, DC, USA: IEEE, 2011, pp. 1747–1752.

[24] W. Liu, J. Lv, B. Yu, W. Shang, and H. Yuan, “Multi-type Road Marking Recognition Using Adaboost Detection and Extreme Learning Machine Classification,” in *2015 IEEE Intelligent Vehicles Symposium (IV)*. Seoul, South Korea: IEEE, 2015, pp. 41–46.

[25] I. M. Chira, A. Chibulcutean, and R. G. Danescu, “Real-Time Detection of Road Markings for Driving Assistance Applications,” in *The 2010 International Conference on Computer Engineering & Systems*. Cairo, Egypt: IEEE, 2010, pp. 158–163.

[26] X. Du and K. K. Tan, “Comprehensive and Practical Vision System for Self-Driving Vehicle Lane-Level Localization,” *IEEE Transactions on Image Processing*, vol. 25, no. 5, pp. 2075–2088, 2016.

[27] F. Poggenhans, M. Schreiber, and C. Stiller, “A Universal Approach to Detect and Classify Road Surface Markings,” in *2015 IEEE 18th International Conference on Intelligent Transportation Systems*. Las Palmas, Spain: IEEE, 2015, pp. 1915–1921.

[28] M. Sukhwani, S. Singh, A. Goyal, A. Behl, P. Mohapatra, B. Bharti, and C. Jawahar, “Monocular Vision based Road Marking Recognition for Driver Assistance and Safety,” in *2014 IEEE International Conference on Vehicular Electronics and Safety*. Hyderabad, India: IEEE, 2014, pp. 11–16.

[29] J. Greenhalgh and M. Mirmehdi, “Automatic Detection and Recognition of Symbols and Text on the Road Surface,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 9493, pp. 124–140, 2015.

[30] O. Bailo, S. Lee, F. Rameau, J. S. Yoon, and I. S. Kweon, “Robust Road Marking Detection and Recognition Using Density-Based Grouping and Machine Learning Techniques,” in *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*. Santa Rosa, CA, USA: IEEE, 2017, pp. 760–768.

[31] T. Chen, Z. Chen, Q. Shi, and X. Huang, “Road Marking Detection and Classification Using Machine Learning Algorithms,” in *2015 IEEE Intelligent Vehicles Symposium (IV)*. Seoul, South Korea: IEEE, 2015, pp. 617–621.

[32] P. Amayo, T. Bruls, and P. Newman, “Semantic Classification of Road Markings from Geometric Primitives,” *International Conference on Intelligent Transportation Systems (ITSC)*, 2018.

[33] R. Danescu and S. Nedevschi, “Detection and classification of painted road objects for intersection assistance applications,” in *13th International IEEE Conference on Intelligent Transportation Systems*. Funchal, Portugal: IEEE, 2010, pp. 433–438.

[34] T.-H. Chan, K. Jia, S. Gao, J. Lu, Z. Zeng, and Y. Ma, “PCANet: A Simple Deep Learning Baseline for Image Classification?” *IEEE Transactions on Image Processing*, vol. 24, no. 12, pp. 5017–5032, 2015.

[35] H. Vokhidov, H. G. Hong, J. K. Kang, T. M. Hoang, and K. R. Park, “Recognition of Damaged Arrow-Road Markings by Visible Light Camera Sensor Based on Convolutional Neural Network,” *Sensors*, vol. 16, no. 12, 2016.

[36] T. Hoang, P. Nguyen, N. Truong, Y. Lee, and K. Park, “Deep RetinaNet-Based Detection and Classification of Road Markings by Visible Light Camera Sensors,” *Sensors*, vol. 19, no. 2, 2019.

[37] Y. Lee, J. Lee, Y. Hong, Y. Ko, and M. Jeon, “Unconstrained Road Marking Recognition with Generative Adversarial Networks,” in *2019 IEEE Intelligent Vehicles Symposium (IV)*. Paris, France, France: IEEE, 2019, pp. 1414–1419.

[38] L. Yann, B. Leon, B. Yoshua, and H. Patrick, “Gradient-Based Learning Applied to Document Recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.

[39] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, “Vision meets robotics: The KITTI dataset,” *International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.

[40] J. Fritsch, T. Kühnl, and A. Geiger, “A New Performance Measure and Evaluation Benchmark for Road Detection Algorithms,” in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*. The Hague, Netherlands: IEEE, 2013, pp. 1693–1700.

[41] J. L. Blanco, F. A. Moreno, and J. Gonzalez, “A collection of outdoor robotic datasets with centimeter-accuracy ground truth,” *Autonomous Robots*, vol. 27, no. 4, pp. 327–351, 2009.

[42] J.-L. Blanco-Claraco, F.-Á. Moreno-Dueñas, and J. González-Jiménez, “The Málaga urban dataset: High-rate stereo and LiDAR in a realistic urban scenario,” *International Journal of Robotics Research*, vol. 33, no. 2, pp. 207–214, 2014.

[43] Y. Kang, H. Yin, and C. Berger, “Test Your Self-Driving Algorithm: An Overview of Publicly Available Driving Datasets and Virtual Testing Environments,” *IEEE Transactions on Intelligent Vehicles*, vol. 4, no. 2, pp. 171–185, 2019.

[44] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, “The Cityscapes Dataset for Semantic Urban Scene Understanding,” in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, 2016, pp. 3213–3223.

[45] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, “1 year, 1000 km: The Oxford RobotCar dataset,” *International Journal of Robotics Research*, vol. 36, no. 1, pp. 3–15, 2017.

[46] F. Yu, H. Chen, X. Wang, W. Xian, Y. Chen, F. Liu, V. Madhavan, and T. Darrell, “BDD100K: A Diverse Driving Dataset for Heterogeneous Multitask Learning,” *arXiv:1805.04687v2 [cs.CV]*, 2020.

[47] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello, “ENet: A Deep Neural Network Architecture for Real-Time Semantic Segmentation,” *arXiv:1606.02147 [cs.CV]*, pp. 1–10, 2016.

[48] H. Zhao, X. Qi, X. Shen, J. Shi, and J. Jia, “ICNet for Real-Time Semantic Segmentation on High-Resolution Images,” in *European Conference on Computer Vision (ECCV) 2018*. Munich, Germany: Springer, Cham, 2018, pp. 418–434.

[49] T. Bruls, W. Maddern, A. A. Morye, and P. Newman, “Mark Yourself : Road Marking Segmentation via Weakly-Supervised Annotations from Multimodal Data,” *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1863–1870, 2018.

[50] J. Janai, F. Güney, A. Behl, and A. Geiger, “Computer Vision for Autonomous Vehicles: Problems, Datasets and State-of-the-Art,” *arXiv:1704.05519 [cs.CV]*, 2017.

[51] Virginia Department of Transportation, “VDOT 2016 Road and Bridge Standards Section 1300 - Traffic Control Devices,” Virginia Department of Transportation, Tech. Rep. September, 2016.

[52] NVIDIA, “GeForce GTX 10710 & 10710 Ti.” [Online]. Available: https://www.nvidia.com/en-in/geforce/products/10series/geforce-gtx-1070/

[53] ——, “GeForce GTX TITAN X.” [Online]. Available: https://www.geforce.com/hardware/desktop-gpus/geforce-gtx-titan-x/specifications

[54] ——, “NVIDIA DRIVE AGX Developer Kit.” [Online]. Available: https://developer.nvidia.com/drive/drive-agx