

# Can we unify monocular detectors for autonomous driving by using the pixel-wise semantic segmentation of CNNs?

Eduardo Romera, Luis M. Bergasa, Roberto Arroyo

**Abstract**—Autonomous driving is a challenging topic that requires complex solutions in perception tasks such as recognition of road, lanes, traffic signs or lights, vehicles and pedestrians. Through years of research, computer vision has grown capable of tackling these tasks with monocular detectors that can provide remarkable detection rates with relatively low processing times. However, the recent appearance of Convolutional Neural Networks (CNNs) has revolutionized the computer vision field and has made possible approaches to perform full pixel-wise semantic segmentation in times close to real time (even on hardware that can be carried on a vehicle). In this paper, we propose to use full image segmentation as an approach to simplify and unify most of the detection tasks required in the perception module of an autonomous vehicle, analyzing major concerns such as computation time and detection performance.

## I. INTRODUCTION

Autonomous driving is a topic with enormous challenges that has recently gathered the interest of teams worldwide in the vehicle industry and research community. Perception is one of its most challenging aspects and requires fusion from multiple sensors to guarantee security by means of redundancies. However, the cheap price and portability of the cameras, the recent advances in computer vision, and the wide possibilities of a high-dimensional signal such as the image (which has not been fully exploited yet) have encouraged the research community to put more efforts on vision-related solutions. One recent example is Mercedes Bertha [1], which supposed one of the first autonomous demonstrations on both interurban and urban environments. Researchers managed to avoid the use of LiDAR by developing a complex perception module that highly relies on a stereo camera, RADAR and multiple monocular detectors that are used to cover the deficits of stereo in the tasks of pedestrian, vehicle, traffic light and lane markings detection.

Traditional approaches for monocular object detection were normally based on classification (e.g. SVM, Adaboost) on top of handcrafted features (e.g. Haar, LBP, SIFT, HOG). The recent appearance of Convolutional Neural Networks (CNNs) has revolutionized the computer vision field due to the impressive learning capabilities that have been demonstrated in complex visual tasks like recognition. The work in [2] presented an approach to combine regions with CNN features (R-CNN) for object detection, that overpassed in

\*This work was supported by the University of Alcalá (UAH) through a FPI grant, the Spanish MINECO through the project SmartElderlyCar (TRA2015-70501-C2-1-R) and the Community of Madrid through the project RoboCity2030 III-CM project (S2013/MIT-2748).

The authors are with the Department of Electronics, UAH, Alcalá de Henares, Madrid, Spain. e-mail: {eduardo.romera, bergasa, roberto.arroyo}@depeca.uah.es

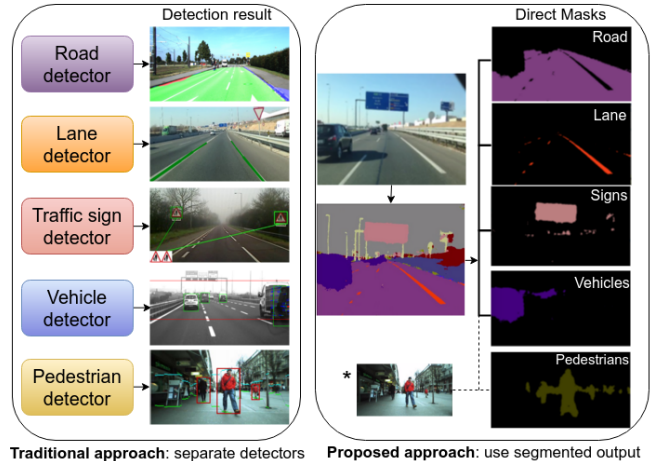


Fig. 1: Two approaches of perception module for autonomous vehicles. \*A different example image was added for pedestrians (due to lack in the first input image) but these are also segmented by the same system.

more than 30% previous best results in PASCAL VOC [3]. More recently, [4] proposed to adapt state-of-the-art CNNs by converting them to Fully Convolutional Networks (FCN) that can be trained end-to-end to perform pixel-wise semantic segmentation. However, an issue with this adaptation is the large size of the receptive fields in the convolutional filters of these networks, which results in coarse outputs with poorly delineated objects. Works like [5] and [6] proposed to solve these problems by combining the network with Conditional Random Fields. While these approaches produce remarkable segmentation, their processing times are a bit far from real-time by ranging from 0.5 to 4 seconds per image on high-end GPUs. However, very recent works like [7] (SegNet) have proposed to combine convolution and deconvolution layers in a new architecture that is optimal for road scene understanding, as it excels in modeling appearance, shape and spatial-relationships (context), producing similar segmentation accuracy with much lower processing times.

These advances have made possible the use of full scene segmentation in time-critical applications like autonomous driving. In this paper, we study how to cover the tasks of the perception module of a vehicle by unifying detection tasks by means of the output of a pixel-wise semantic segmentation network (see Fig. 1). The main difference of such approach is that, instead of performing classification on pre-selected region candidates (traditional methods), the whole image is classified to use this output as base to higher-level understanding algorithms required for vehicle navigation (e.g. collision avoidance). In the following sections we analyze the advantages and challenges of the proposed approach as compared to traditional monocular detection approaches.

## II. DISCUSSION ON GENERAL CONCERNS

We analyze the general advantages and disadvantages of unifying detectors by using pixel-wise segmentation:

1) **Computational cost and latency:** Until recently, full pixel segmentation was not usable in terms of efficiency. However, recent works such as SegNet [7] allow to compute decent segmentation (see Fig. 2) in less than 100ms on a single GPU. If the task is only to detect one single category, then any monocular detector can outperform this performance. However, in autonomous vehicles, it is presumed that we are interested in detecting everything that surrounds the vehicle. In this case, running all detector jointly scales computation much further than 100ms (see Table I), so the unification in this case is advantageous. A critical point may be the latency: fast moving objects are response-time critical (e.g. pedestrians in street scenarios), and the unification detracts these cases as the latency is the same for all detections.

2) **Implementation complexity:** One of the main advantages of the unification is that it simplifies the problem. In the case of independent detectors, each approach is based on a different model, algorithm and training, which implies high complexity for implementation and for developing algorithms on top of this result. Additionally, optimization is not always straightforward (e.g. parallelization or GPU implementation). However, in the case of using the segmented output of a CNN, every implementation challenge unifies within the same module. Recent deep learning libraries like Caffe [8] have simplified development (shared code), optimization (direct GPU implementation) and even have reinforced the community implication on sharing pre-trained models (Model Zoo). Upgrading the system is also straightforward, both in the case of changing network architecture and in the case of re-training with new data (easy to fine-tune parameters instead of having to re-train multiple classifiers).

3) **Output quality:** In the case of pixel-wise segmentation, even if the output is not precise, the result might still be of use. For instance, in obstacle detection it is preferred to know that there is a pedestrian in some part of the image (even if the shape is not accurate) to not having anything detected at all. Additionally, for cases in which shape is important (i.e. traffic signs) the segmented output may be used as a "ROI detector" to scan object candidates and apply a simple classifier in that area (with reduced computational cost). Moreover, having the full pixel segmentation provides additional advantages which are not only related to detection, but also to higher-level applications such as semantic reconstruction of the environment.

4) **Data availability and training:** In the case of traditional detectors, each one requires a set of positives and negatives to train each specific class. With the recent progress on CNNs, pixel-annotated datasets are becoming popular and more detailed, so training end-to-end to perform pixel segmentation has become feasible and a good practice to jointly optimize the network weights. PASCAL VOC 2012 [3] is one of the first with fully labeled images (indoor and outdoor). More recently, datasets oriented towards vehicle environments have emerged, like CamVID [9] and Cityscapes [10].




Fig. 2: Example cases of pixel-wise segmentation performed by SegNet on real road scenarios: daytime (top), dusk (middle) and nighttime (bottom)

5) **Tracking:** While traditional approaches like [2] firstly extract region candidates for posterior classification, in the proposed approach no region candidates are extracted before classification (segmentation) occurs. Therefore, if tracking of individual objects is desired, then the object candidates must be extracted a posteriori. Some possibilities are to extract object candidates by applying shape constraints or enclosing those with same motion (e.g. using optical flow). Additionally, the unified knowledge about the whole road environment allows for better application of inter-class relations to improve tracking. For instance, some vehicle detectors like [11] rely on lane detection to track according to geometrical constraints and to filter impossible vehicle locations. By having full segmentation, these relations are easier to apply (see Fig. 3 for an example). In fact, the spatial-relationship between different classes, such as the difference between road and side-walk, is implicitly learned by the CNN.

6) **Occlusions:** Works like [12] reflect that occlusions suppose a big impact on detection performance on traditional detectors, even for partial occlusions, and it surprisingly affects similarly to both holistic and part-based detectors. By using the pixel-wise segmented output this effect is reduced, as we do not need region candidates that enclose the whole object, but even a small object region can produce enough feature activations to segment correctly the class on those pixels (unless the occlusion occurs in the parts with the most representative features of the specific class).

7) **Challenging lighting conditions:** While detection decay is severe in any traditional detector at challenging conditions (night or adverse weather), the accuracy in the segmentation produced by a CNN remains still reasonable (see Fig. 2 for dusk and nighttime examples). At night, areas covered by the vehicle headlights (road and close signs) are correct, while distant objects might be extremely noisy and may require combination with additional detection methods.




Fig. 3: Example application of the segmented output (direct extraction of lane and vehicle positions). Left: input and segmented output. Mid: bird's view (IPM) of input. Right: bird's view of segmented lanes and vehicles.

### III. DISCUSSION ON SPECIFIC DETECTORS

We discuss the advantages and drawbacks on specific detection concerns in five of the main categories an autonomous system should perceive. Table I presents a quantitative comparison that summarizes the reviewed methods.

1) **Road:** Detecting the road is necessary for delimiting the driving space. Bertha [1] covered this part with stereo vision by treating ground plane as the road. However, stereo vision has its drawbacks, as its range may not be sufficient for high speeds (i.e. motorways) and the ground plane may not necessarily be asphalt road in some environments (e.g. grass, low sidewalks or water). Some works have treated road detection with monocular cameras like [14], that achieved near real-time pixel-wise output: around 50ms per frame (4 cores@2.8GHz) with a precisions between 77%-88% and recall rates between 83%-96%. However, the computation impact is still high considering that all the cores of a high-end CPU are being used to only detect the road. Therefore, in this case there is no specific reason for using an independent detector for road instead of the segmented output of a CNN, which in fact provides better detection performance (94.5% pixel accuracy in the road class, as seen in Table II).

2) **Lane markings:** Detecting own lane is required for avoiding invasion of other lanes which may be occupied by other participants or that may be reserved for the opposite direction. Current vehicles like [1] rely the knowledge of lane spatial properties to a pre-designed map. There has been active research on lane detection algorithms with a wide variety of lane representations (e.g. fixed-width line pairs, spline ribbon), detection and tracking techniques (from Hough transform to probabilistic fitting and Kalman filtering). A simple approach for smartphones that used a clothoidal model on a EKF (based on points extracted from an adaptive Canny algorithm) is presented in [13]. A robust probabilistic approach suited for all kind of challenging scenarios (worn lanes and different kinds of splitting lanes) was presented in [15]. Computation cost in these detectors is low. [13] can run at 30 FPS on a single core on iPhone 5, and [15] in less than 25ms per frame on an Intel Core@1.83GHz (2008), even faster in recent processors. Because of this, using full segmentation only to extract lanes could be disproportionate, but it makes sense if the segmentation is used

for several tasks as in the proposed approach. From the segmented output, it is straightforward to obtain a lane model by using an EKF (with clothoidal model for curvature) on the segmented lane-markings observations from a bird's view perspective (like in Fig. 3). An additional concern are the roads with no visible markings. Bertha used a curb detector to consider them a boundary line (mainly for suburban areas). In the case of full segmentation, approaches like SegNet allow to separate classes like sidewalk and fence, so in combination with the road detection it would be simple to detect them as a lane boundary if there is no lane marking.

3) **Traffic signs and lights:** Detecting traffic signs and lights is needed for understanding the main navigation rules. However, this is a challenging task due to the high sign variability, different regulations and the spatial implications (e.g. which signs apply to which lanes). Reader may refer to [16] for a study on traffic sign recognition methods. Current autonomous vehicles rely all this information to highly detailed maps and they only detect dynamic properties like the states of traffic lights. The main problem of an approach based on full segmentation is that no sign type classification is performed, so the segmentation could help in locating possible sign/light candidates in the image, but still an additional algorithm would be required to extract information about the properties. Therefore, there is no direct computation saving in the classification on the proposed approach (only in the extraction of sign candidates). However, applying a simple classification to infer the sign type or the state of a traffic light is relatively simple if the region candidate is provided.

4) **Vehicles:** It is essential to detect other traffic participants that share the road (i.e. cars, trucks, buses, bicycles and motorbikes) to avoid collision and also to understand their position and intentions in order to anticipate and respect their priority (e.g. not invade their driving lane). Recently, [11] introduced a detection and tracking framework for smartphones by applying geometrical constraints and using multiple regions to reduce scanned amount of pixels. However, this results in strong decay in detection from 70 meters onwards, which may not be usable for autonomous driving. Other approaches like [17] propose a two-part detection technique using an active learning framework and symmetry-based iterative analysis, providing state-of-the-art results in around 20FPS on 4 cores (i7). The problem of these approaches is that they are normally trained and tested only on cars (and trucks at most). Different vehicles have varied shapes, which results in an important detection decay even between similar vehicles (see [11] for car vs trucks), and other kinds like motorbikes or bicycles will not be detected at all. In such case, separate classifiers trained for each vehicle category would be desired, but it is not realistic due to computational constraints. Due to these reasons, the proposed approach based on full segmentation would be advantageous as it differs and detects varied kinds of vehicles (as required by an autonomous vehicle) with remarkable detection rates (see Table II). Distinguishing between type of vehicles also adds the possibility to use different prediction models (e.g. a car

TABLE I: Summarized quantitative comparison between monocular detectors (left) and pixel-wise segmentation networks (right). It should be considered only as an overview, as direct performance comparison is not possible due to different evaluation metrics on each task. \*Bergasa [13] and Romera [11] are smartphone approaches, but times are measured in computer by using an iPhone 6 simulator. \*\*IoU means Intersect over Union = TP/(TP+FP+FN), metric introduced in [3] to measure pixel classification accuracy.

Task	Method	Time per frame [environment, resolution]	Performance (different metrics)	Network	Time per frame (forward pass on 500x500)	Performance IoU** (Top: PASCAL VOC [3], Bot: Cityscapes [10])
Road	Passani [14]	50 ms [4 cores i7@2.1Ghz, 375x1242]	Precision: 78-89% Recall: 82-93%	FCN-8s [4]	500 ms (published time)	62.2% class avg
Lane	Bergasa [13]	5 ms [4 cores i7@2.2Ghz*, 640x480]	n/a			65.3% class avg 85.7% category avg
	Kim [15]	25 ms [Intel Core@1.83Ghz, 70x250]	n/a	Deeplab [5]	4000 ms (published time)	58% class avg
Signs	(various) [16]	N/A	72 to 99% mean detection rate			64.8% class avg 81.3% category avg
Vehicles	Romera [11]	16 ms [4 cores i7@2.2Ghz*, 640x480]	Prec: +90% Recall (<60m): 80 to 95%	CRFasRNN [6]	700 ms (published time) ~2 sec. (tested Nvidia Titan X)	69.6% class avg
	Satzkoda [17]	50 ms (4 cores i7, 1024x768)	Prec: 94% Recall: 97%			62.5% class avg 82.7% category avg
Pedestrians	Dollar [12]	150-350 ms [2012 CPU n/a, 640x480]	57% log-average miss rate (on clearly visible pedestrians)	SegNet [7]	60 ms (published time) ~100 ms (tested Nvidia Titan X)	59.1% class avg 79.1% category avg

TABLE II: Pixel accuracy per class (%) of SegNet trained on a multiple dataset (3.5K images) evaluated on CamVid [9] dataset (11 classes)

SegNet (3.5K)	Build.	Tree	Sky	Car	Sign	Road	Pedes.	Fence	Pole	Sidew.	Bike	Class Avg.	Global Avg.	Mean IU
73.9	90.6	90.1	86.4	69.8	94.5	86.8	67.9	74.0	94.7	52.9	80.1	86.7	60.4	

moves differently to a bike). Additionally, the simultaneous knowledge of the road and lanes promotes easier tracking (known vehicle position with respect to lanes, as in Fig. 3) or filtering of impossible locations (e.g. vegetation, fence).

5) **Pedestrians:** Robust pedestrian detection is a key aspect in autonomous driving, specially in urban scenarios. The survey [12] presents an extensive evaluation of state-of-the-art pedestrian detectors in a unified framework. Nearly all modern detectors employ as features some form of gradient histograms, with SVM and boosted classifiers as the most popular learning algorithms. A widely known detector, with one of the best detection performances and fastest computation is FPDW [12], processed between 2.7 and 6.5 FPS on CPU (640x480). Although additional improvements can enhance this result by speeding the pre-selection of region candidates or the detection itself, this task still supposes heavy computation load to the framework and the detection rates may not be robust enough for autonomous driving (miss ratio is still significant and there is a heavy detection decay with occlusions and candidates of small pixel sizes). Therefore, relying the pedestrian detection task on the pixel-wise segmentation (proposed approach) would be advantageous in both computation (see Table I) and detection performance (86.8% pixel accuracy as seen in Table II).

#### IV. CONCLUSION AND FUTURE WORK

This work supposes a prior study to design a perception approach for autonomous vehicles fully based in segmentation performed by CNNs, which will be reviewed in future works. The aim is to simplify things in terms of implementation and to solve open challenges of traditional detectors such as the high combined computation cost and detection decay in cases like occlusion. However, such an approach also adds challenges of its own (achieved segmentation is not perfect), but end-to-end CNNs are a recent topic and it still has room for future improvements in both efficiency and accuracy, where new datasets and network architectures will play an essential role.

#### REFERENCES

- [1] J. Ziegler, P. Bender, M. Schreiber, H. Lategahn, T. Strauss, C. Stiller, T. Dang, U. Franke, N. Appenrodt, C. G. Keller *et al.*, "Making
- bertha drive - an autonomous journey on a historic route," *IEEE Intel. Transportation Systems Magazine*, vol. 6, no. 2, pp. 8–20, 2014.
- [2] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *IEEE Conf. on Computer Vision and Pattern Recog. (CVPR)*, 2014, pp. 580–587.
- [3] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The Pascal visual object classes (VOC) challenge," *International Journal of Computer Vision*, vol. 88, no. 2, pp. 303–338, 2010.
- [4] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *IEEE Conf. on Computer Vision and Pattern Recog. (CVPR)*, 2015, pp. 3431–3440.
- [5] L.-C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, "Semantic image segmentation with deep convolutional nets and fully connected crfs," *arXiv preprint arXiv:1412.7062*, 2014.
- [6] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. Torr, "Conditional random fields as recurrent neural networks," in *IEEE International Conf. on Computer Vision (ICCV)*, 2015, pp. 1529–1537.
- [7] V. Badrinarayanan, A. Handa, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for robust semantic pixel-wise labelling," *arXiv preprint arXiv:1505.07293*, 2015.
- [8] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell, "Caffe: Convolutional architecture for fast feature embedding," in *ACM International Conf. on Multimedia*, 2014, pp. 675–678.
- [9] G. J. Brostow, J. Fauqueur, and R. Cipolla, "Semantic object classes in video: A high-definition ground truth database," *Pattern Recognition Letters*, vol. 30, no. 2, pp. 88–97, 2009.
- [10] 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 *IEEE Conf. on Computer Vision and Pattern Recog. (CVPR)*, 2016.
- [11] E. Romera, L. M. Bergasa, and R. Arroyo, "A real-time multi-scale vehicle detection and tracking approach for smartphones," in *IEEE Intel. Transportation Systems Conf. (ITSC)*, 2015, pp. 1298–1303.
- [12] P. Dollar, C. Wojek, B. Schiele, and P. Perona, "Pedestrian detection: An evaluation of the state of the art," *IEEE Trans. on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 34, no. 4, pp. 743–761, 2012.
- [13] L. M. Bergasa, D. Almería, J. Almazán, J. J. Yebes, and R. Arroyo, "DriveSafe: an app for alerting inattentive drivers and scoring driving behaviors," in *IEEE Intel. Vehicles Symp. (IV)*, 2014, pp. 240–245.
- [14] M. Passani, J. J. Yebes, and L. M. Bergasa, "Fast pixelwise road inference based on uniformly reweighted belief propagation," in *Intelligent Vehicles Symposium (IV)*, 2015 IEEE. IEEE, 2015, pp. 519–524.
- [15] Z. Kim, "Robust lane detection and tracking in challenging scenarios," *IEEE Trans. on Intel. Transportation Systems*, vol. 9, pp. 16–26, 2008.
- [16] A. Mogelmoose, M. M. Trivedi, and T. B. Moeslund, "Vision-based traffic sign detection and analysis for intelligent driver assistance systems: Perspectives and survey," *IEEE Trans. on Intel. Transportation Systems*, vol. 13, no. 4, pp. 1484–1497, 2012.
- [17] R. K. Satzoda and M. M. Trivedi, "Multipart vehicle detection using symmetry-derived analysis and active learning," *IEEE Trans. on Intel. Transportation Systems*, vol. 17, no. 4, pp. 926–937, 2016.