https://doi.org/10.2352/ISSN.2470-1173.2021.17.AVM-210

RoadEdgeNet: Road Edge Detection System Using Surround View Camera Images

***Ashok Dahal* [1], *Eric Golab* [2], *Rajender Garlapati* [1], *Varun Ravi Kumar* [3] *and Senthil Yogamani* [4]**
[1] ***Valeo DVS Troy, MI USA***
[2] ***University of Michigan, Ann Arbor, MI USA***
[3] ***Valeo DAR Kronach, Germany***
[4] ***Valeo Vision Systems, Ireland***

Abstract

Road Edge is defined as the borderline where there is a change from the road surface to the non-road surface. Most of the currently existing solutions for Road Edge Detection use only a single front camera to capture the input image; hence, the system's performance and robustness suffer. Our efficient CNN trained on a very diverse dataset yields more than 98% semantic segmentation for the road surface, which is then used to obtain road edge segments for individual camera images. Afterward, the multi-cameras raw road edges are transformed into world coordinates, and RANSAC curve fitting is used to get the final road edges on both sides of the vehicle for driving assistance. The process of road edge extraction is also very computationally efficient as we can use the same generic road segmentation output, which is computed along with other semantic segmentation for driving assistance and autonomous driving. RoadEdgeNet algorithm is designed for automated driving in series production, and we discuss the various challenges and limitations of the current algorithm.

Introduction

Road edge detection is a significant problem for applications in autonomous driving systems. Such systems must detect the limit of the drivable road surface in multiple types of weather and lighting conditions. As such, this is not a trivial task. Additionally, due to the vehicle's high velocity, the algorithm must process incoming images very rapidly, on the order of less than 100 ms, to be of practical use for automotive systems. Different pipelines exist for road edge boundary detection algorithms. Work in this domain includes approaches that use traditional machine learning algorithms with hand designed features and more advanced techniques that make use of convolutional neural network architectures. Traditional approaches use classifiers with feature vectors constructed from color or histogram data. These are then used to determine probabilistically where the road edge is located. Designing features that adequately represent the road edge is a necessity in such systems. Often, the algorithm's final performance is determined by the skill of the programmer in choosing meaningful features.

Convolutional neural network based approaches typically identify the presence of a class in a particular image. Large neural networks such as AlexNet or VGGNet are examples of this principle [1] [2]. The last layer of the neural network can be removed to approximate where the class is in the picture. However, this is a rough approximation, and correct semantic segmentation must be accomplished using a network of convolution and deconvolution operations. SegNet is an example of this type of network architecture [3]. It is a symmetric network that performs semantic segmentation on an image. The first half of the network performs convolution, ReLU, and downsamples the input image. The second half performs upsampling and deconvolution. SegNet can process a 360 x 480 image in 422.5 ms on NVIDIA Titian GPU with cuDNN v3, achieving a mean intersection over union (mIoU) score of 60.10 on all classes in the CamVid dataset [4]. An even more advanced architecture available is E-net [5]. E-net consists of a network of bottleneck operations. These bottlenecks have been shown to be easier to generalize than traditional convolution operations [6, 7]. Additionally, E-net also uses asymmetric filters of 1x5 and 5x1 to approximate a 5x5 filter's action. This reduces the number of parameters in the network, which reduces training time and makes the network faster to execute on streamed image data [8]. In their work, parametric ReLU (PReLU) was also found to produce better results over ReLU. Both SegNet and E-net perform multi-class semantic segmentation of images. Each network is capable of segmentation of multiple classes in a scene. It is claimed that E-net is capable of semantic segmentation on a 480x320 image in 47 ms with an NVIDIA TXI and 7 ms for an NVIDIA Titian [5]. Such short processing speeds allow for real-time applications.

Road Edge is defined as the borderline where there is a change from the road surface to the non-road surface. Detection of the road edge is crucial for driving assistance systems. It gives the vehicle a higher-level understanding of the local environment (e.g., intersections) and is especially useful when there are no painted lane markers (e.g., rural roads). Most of the currently existing solutions use only a single front camera to capture the input image; hence, the system's performance and robustness suffer [9]. They have also used traditional computer vision algorithms that have recently been shown to have far inferior performance compared to deep learning-based systems. Some related works use multiple cameras to create a virtual bird's eye view image. That is a single perspective-corrected image created by warping and stitching together the individual raw images from multiple cameras [10]. This approach has the advantage of fusing information from multiple cameras, which is handled as a pre-processing step before the computer vision algorithm receives the data. However, this approach is very sensitive to the camera system, intrinsic and extrinsic calibration errors, and it is also very dependent on the resolution of the camera images. These effects reduce the effective measurement range of such systems (assuming 1 MP cam-

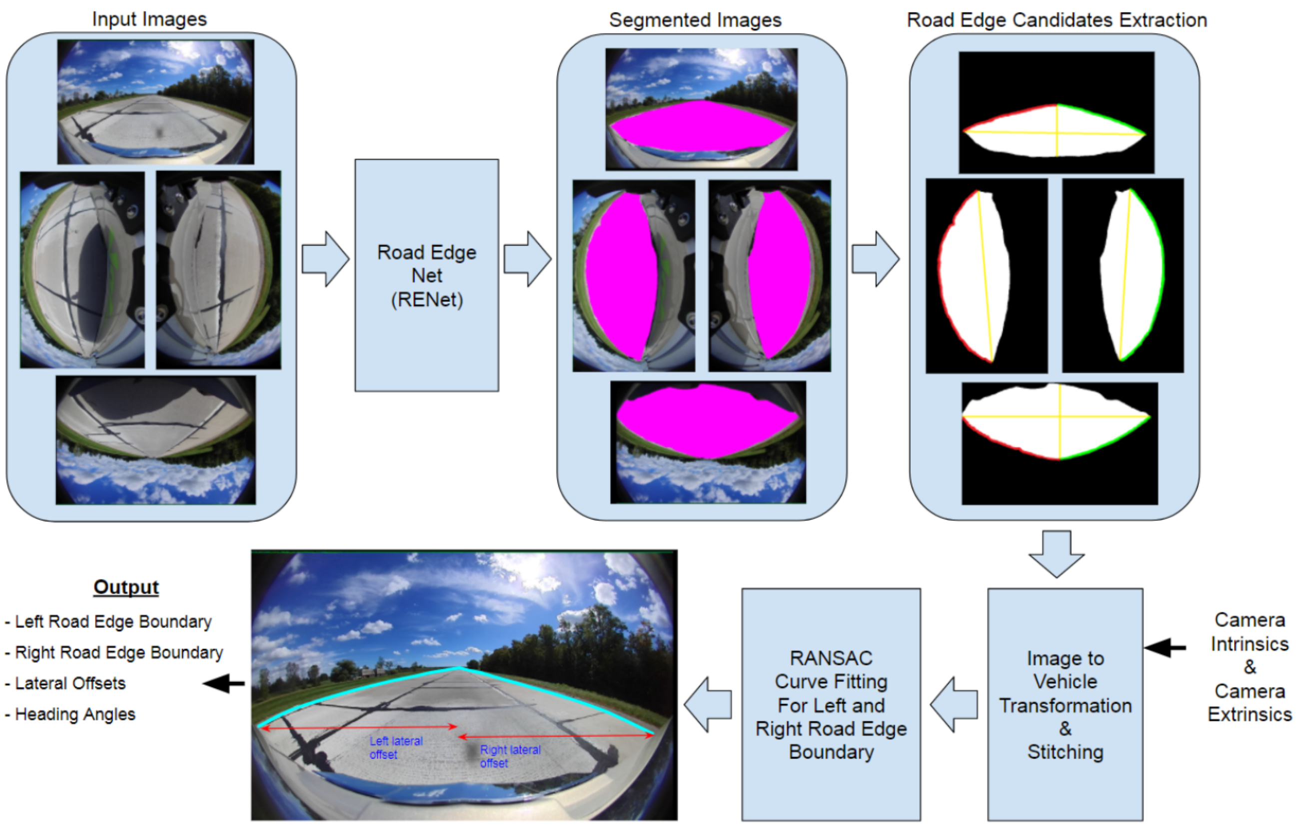

Figure 1: **Overall Road Edge Detection System Architecture**

eras) to approximately 5 meters surrounding the vehicle. Thus, this is not a preferred approach for high-speed applications such as autonomous driving, where measurement ranges of at least 15 meters are required.

Benefits of Our Surround View Approach

Multiple cameras provide an overlapped view of the same scene, which improves the detection robustness from occlusions (other vehicles or objects) and environmental conditions (weather conditions, lens soiling). Our surround-view cameras capture multiple images so that road edges occluded from the front or any other camera can be captured by the remaining cameras. It is not guaranteed that a single camera will detect the road edges for every single frame of the scene; this is because of the weather conditions, time of day, lighting conditions, occlusions, etc. The presence of multiple cameras increases the chance of detection, and accuracy increases significantly.

Relatively there is less literature on surround view cameras but there has been recent progress in various visual perception tasks such as semantic segmentation [11], moving object detection [12], depth estimation [13, 14, 15, 16], re-localisation [17], soiling detection [18, 19, 20] and multi-task learning [21, 22]. To overcome the limitations of the bird's eye view approach, we purpose performing the road edge feature detection on the raw fisheye camera images before projecting detected features into the vehicle coordinate system and fusing to create a combined road edge detection system. This approach allows us to avoid warping artifacts due to the resolution of the camera system. The mapping from camera pixel coordinates to vehicle coordinates (also known as an image to ground transformation) is performed only for detected and filtered contours compared to existing surround-view solutions where all pixels are first transformed. This significantly improves the runtime performance and saves a significant amount of memory. Unlike existing solutions, we use our efficient Convolutional Neural Network-based semantic segmentation to detect the road area, from which we extract candidate road edge segments. This has a significant impact and improves performance even in very challenging scenes. Compared to current CNN networks such as ENet [5] for semantic segmentation, our network 1) has almost 2X fewer network parameters, 2) is 2X faster in terms of runtime at the same time is better or on par with those methods for segmenting road and non-road regions. Our unique candidate road edge extraction method is efficient in terms of run time, as we do not go through all pixels in the segmented image while looking for edges.

System Architecture

The Road Edge Detection system receives four images from the Front, Rear, Mirror Left, and Mirror Right Cameras. These images are passed, at once, as one batch to the Road Edge Convolutional Neural Network (RoadEdgeNet), which produces four segmented images simultaneously. Once the segmentation is complete, the raw candidate edge segments from all four contours are extracted. Each edge segment contains many candidate edge points. These candidate edge points are transformed to the vehicle coordinate system, and similar feature points detected in multiple cameras are stitched together. At this point, there are candidate left and candidate right road edge feature points. These candi-

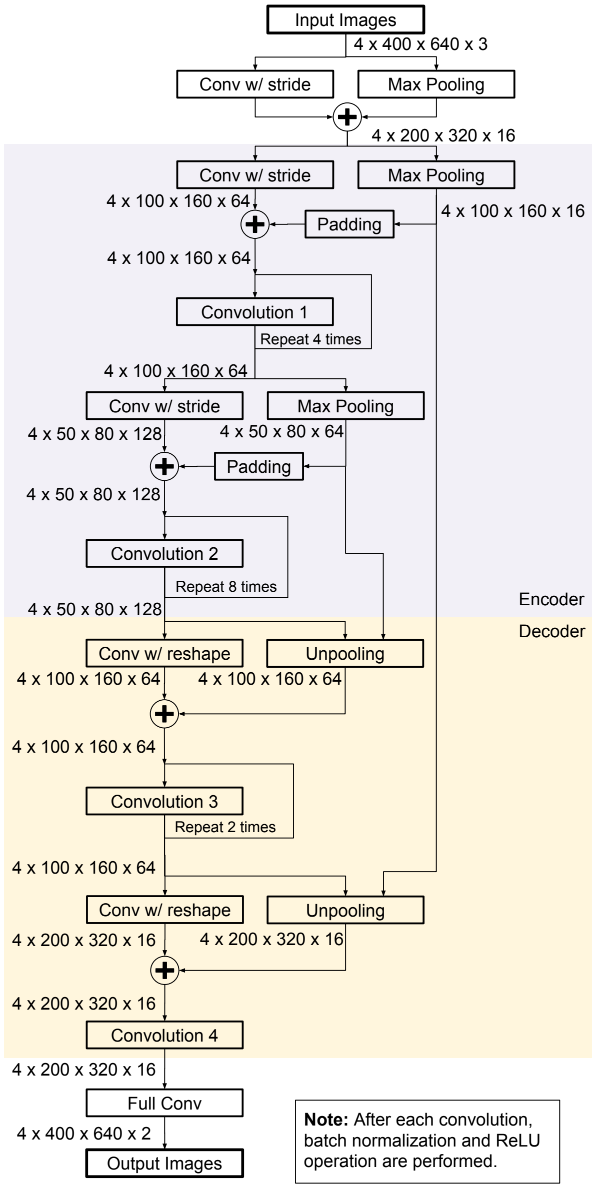

Figure 2: **Our Proposed RoadEdgeNet Architecture**

dates left and right feature points are fitted with the RANSAC curve fitting algorithm to get the final road edge boundaries. The fitted RANSAC boundaries, lateral offsets, and heading angles concerning the left and right boundaries are calculated. This information can then be used for autonomous driving applications, such as to keep the vehicle inside the road in case there are no lane markers present (e.g., rural roads) or to observe the high-level information around urban driving.

RoadEdgeNet is a system that builds upon convolutional neural networks such as SegNet and E-net for semantic segmentation. One significant difference between RoadEdgeNet and the work described above is that RoadEdgeNet only performs semantic segmentation on dual-class images. This simplification allows the network to locate only drivable road surfaces, which allows the network size to be reduced. The overall road edge detection system can be described in the following way. The overall system pipeline is shown in Figure 1. The Road Edge Detection system captures four Front, Rear, Mirror Left, and Mirror Right Cameras images. These images are passed to RoadEdgeNet, which produces four segmented images simultaneously (Figure 1 top right). Our typical frame size is 800x1280 coming from camera sensors. We downsample by half to 400x640 before feeding to the neural network as shown in Figure 2. Initially, the network further reduces the input image size and increases the feature map's depth to 16. The information is provided to the encoder block (shown in purple). It extracts the abstract features followed by a decoder block (shown in yellow) where encoder features are mapped to the pixel level for semantic segmentation. At the end of the network, the image is segmented into two classes based on each pixel's classification. For each convolution operation, 1x1 projection is performed before the primary convolution to reduce the feature map's depth. This operation's benefit is 1) reducing the number of parameters, 2) runtime/speed improvements. The depth of the feature map is restored using another 1x1 projection at the end. In between each convolution, batch normalization and ReLU (activation function) is also performed.

Extracting Road Edges from Semantic Segmentation Output

After the semantic segmentation of the road and the non-road surface, we utilize the road semantic segmentation results to extract preliminary road edge candidate points from each camera described in Figure 4. Each edge segment contains many candidate edge points. These candidate edge points are transformed to the vehicle coordinate system, and similar feature points detected in multiple cameras are stitched together, as shown in Figure 3. To transform the points from image to vehicle coordinates, camera calibration parameters are required, which are called extrinsic and intrinsic parameters. The world points are transformed into camera coordinates using the extrinsic parameters. The camera coordinates are mapped into the image plane using the intrinsic parameters. Once image coordinate points are mapped to the world/vehicle coordinates system, each point's distance concerning the vehicle can be found. This is crucial for stitching similar road segments and finding which points belong to left or right road edge boundaries. The transformation is done as follows:

$$W[xy1] = [XYZ1]P \tag{1}$$

where $P = cameramatrix = [Rt]K$; W is the scale factor and in our case, $W = 1$ (no scaling); $[xy1] = imagepoints$; $[XYZ1] = worldpoints$; R is the rotation matrix; t is the translation matrix and K is the camera intrinsic matrix.

Once the edge points transformation is done, RANSAC (an iterative method) is used to estimate parameters of a mathematical model (curve/line) from a set of observed data that contains outliers. Once stitching is completed, we will have a set of points for the left and right road edge boundaries. Some of the points will be outliers, but they will be filtered out using the RANSAC fitting method, as seen in Figure 3 (right). Once the final left and right road edges are determined, they are overlayed on the front camera image to show the estimated road edges. The proposed algorithm can also be extended to be used with other sensors like ultrasonics [23], Lidar [24, 25] and radar.

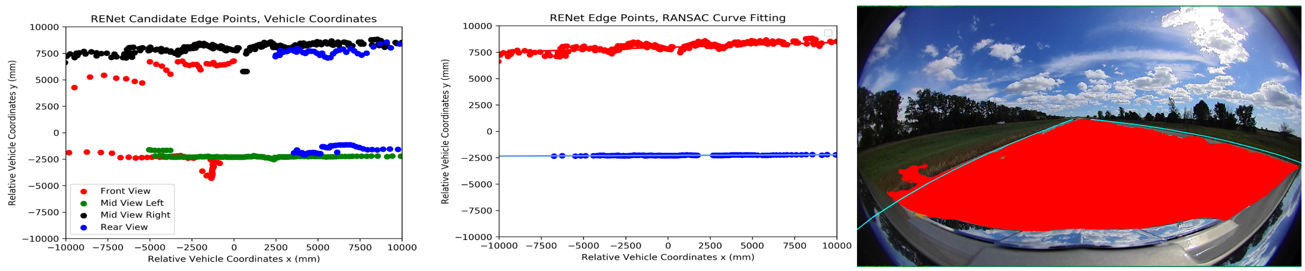

Figure 3: **Road Edges candidate points** from 4 cameras (left), road edge points from RANSAC curve fitting (middle) and Final Road Edges Overlay on the Front Camera (right)

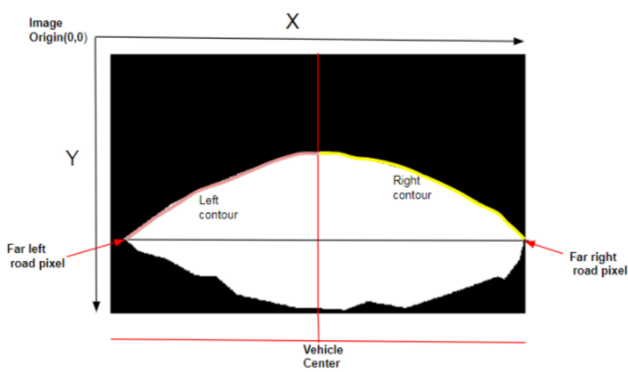

Figure 4: **Obtaining candidate left and right road edge points.** Note that this figure illustrates the Front View camera image. The same logic is applied for the Rear View image. For mirror cameras, we will perform steps 1 - 3 only as we would only have either left or right side edge points on mirror cameras. Step-1: Get the far left road pixel (x1, y1) and far-right road pixel (x2, y2) from the segmented binary image. Step-2: Scan from left to right for each row. If a road pixel is reached, store the point and go to the next row while skipping other pixels in that row. Step-3: Repeat step 2 for each row until y1 is reached and skip everything below y1. Step-4: Scan from right to left for each row. If a road pixel is reached, store the point and go to the next row skipping all other pixels in that row. Step-5: Repeat step 4 for each row until y2 is reached and skip everything below y2. Step-6: Transform the points into vehicle coordinates with the Image to World Transformation for further processing.

Metrics	E-Net			RoadEdgeNet		
	Valeo	CamVid	Synthia	Valeo	CamVid	Synthia
Avg. class accuracy	0.97	0.98	0.98	0.97	0.98	0.98
Total accuracy	0.97	0.99	0.98	0.97	0.98	0.98
Mean IoU	0.94	0.97	0.97	0.94	0.97	0.96
Class 0	0.97	0.99	0.99	0.98	0.99	0.99
Class 1	0.96	0.98	0.98	0.96	0.98	0.97
Total Precision	0.97	0.99	0.99	0.98	0.99	0.99
Total Recall	0.98	0.99	0.99	0.98	0.99	0.98
Total F1 score	0.98	0.99	0.99	0.98	0.99	0.98
Total MCC score	0.94	0.97	0.97	0.94	0.97	0.96
Runtime(sec/step)	0.50	0.38	0.23	0.15	0.14	0.13

Table 1: **Road Semantic Segmentation** results on different datasets on the E-Net and RoadEdgeNet.

Experiments and Results

RoadEdgeNet is trained on a diverse set of images collected by the authors. These images include multiple weather and lighting conditions, including snow, rain, cloudy sky, dawn, and dusk, as well as night. The collection of images represents all four surround-view cameras giving different perspectives of the scene. The total dataset consisted of 682 images split into 385 for training, 90 for validation, and 293 for testing using sampling techniques described in [26]. The vehicle and sensor setup used is same as in our WoodScape dataset [27]. Unlike other datasets such as the CamVid dataset, [28] and the SYNTHIA dataset [29] Valeo's dataset only contains two semantic classes, road, and non-road. The table 1 shows the semantic segmentation results between ENet and RoadEdgeNet.

Both networks are also trained and evaluated on the CamVid dataset consisting of 701 images with 420 train, 70 validation, and 211 test images. Additionally, they are also trained and evaluated on a subset of the SYNTHIA dataset consisting of 405 images with 209 train, 37 validation, and 159 test images. All experiments were performed on an NVIDIA Quadro M1000M. In the Automotive domain, we have minimal computation power for embedded systems, so our goal is to reduce the inference run-time while maintaining high segmentation accuracy. These results illustrate that our Network can achieve that goal as run-time is more than 2X better than ENet, and accuracy is very similar on different measurement metrics as shown in Table 1. Semantic segmentation results perform well enough that the system is appropriate for rural roads with no lane markers. Consistent performance is seen high dynamic range (HDR) environments and low lighting conditions. This can be used for autonomous driving to keep in the vehicle within the road region. As far as network parameters are concerned, RoadEdgeNet has 0.2M vs ENet's 0.37M, which is almost half as many which resulted in faster performance.

Future Works and Challenges

Road edges are present in various textures such as sands, rocks/gravels, dirt, grass, snow/sleet and so on. It is challenging for edge detection especially in these scenarios where border of road surface and various textures can be too deceptive of edge identification. A significant amount of diverse catalog (various textures) based training will be beneficial for achieving good performance on various conditions. As there are no public datasets for road edge detection, it will be beneficial to make our dataset public so others can add on it to make it more diverse. We plan to make our internal Valeo dataset public for this reason. We also plan to improve our dataset by adding more diverse images from various scene catalogs and fine tune our algorithms to better work on those diverse conditions.

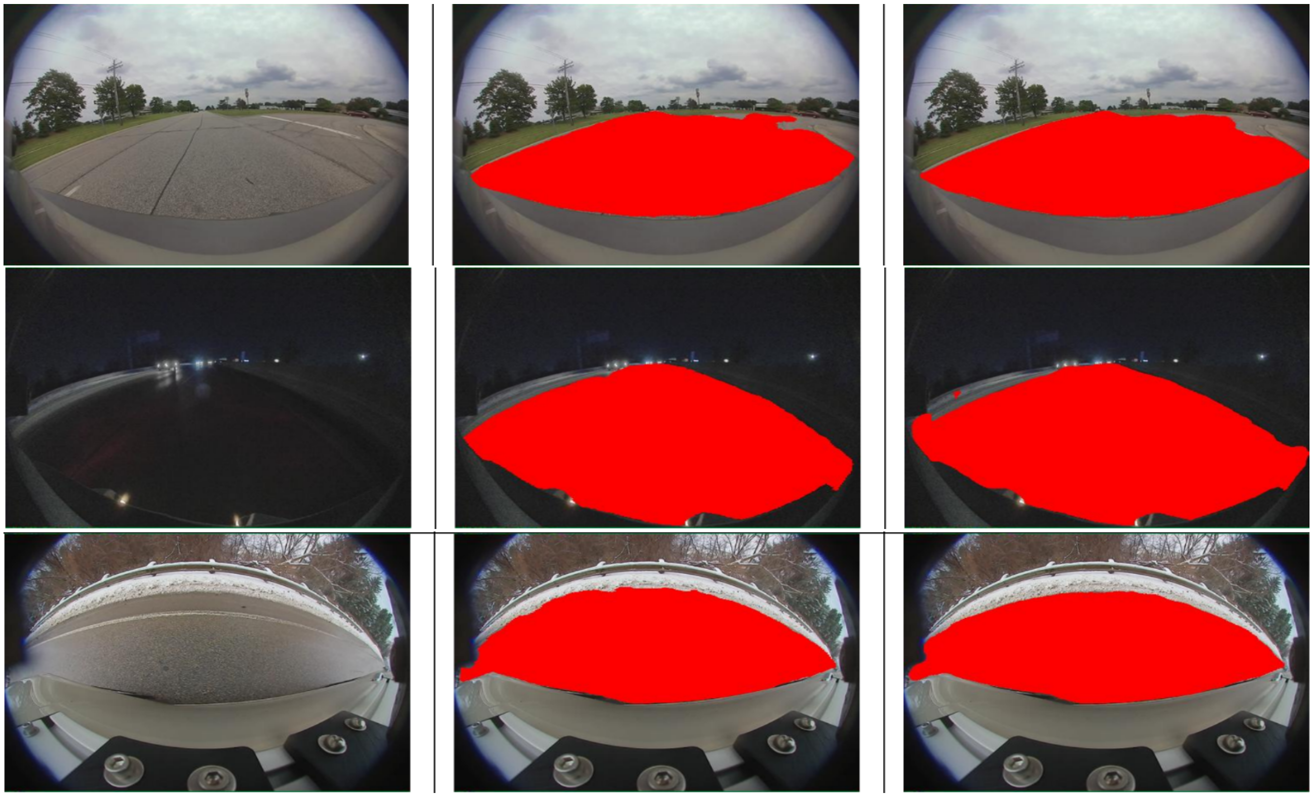

Figure 5: **Semantic Segmentation Results on Valeo Dataset.** Original Image (left), ENet (middle) and RoadEdgeNet (right).

Figure 6: **Semantic Segmentation Results on CamVid Dataset.** Original Image (left), ENet (middle) and RoadEdgeNet (right).

References

[1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *ImageNet Classification with Deep Convolutional Neural Networks*. arXiv, 2012.

[2] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv*, pp. 1–14, 2015.

[3] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation." 2015.

[4] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder- decoder architecture for image segmentation," *arXiv*, pp. 1–14, 2016.

[5] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciel, "Enet: A deep neural network architecture for real-time semantic segmentation," *arXiv*, pp. 1–10, 2016.

[6] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *arXiv*, pp. 1–12, 2015.

[7] K. He, X. Zhang, S. Ren, and J. Sun, "Delving deep into rectifiers: Surpassing human- level performance on imagenet classification,"

Figure 7: **Semantic Segmentation Results on SYNTHIA Dataset.** Original Image (left), Enet (middle) and RoadEdgeNet (right).

arXiv, pp. 1–11, 2015.

[8] J. Jin, A. Dundar, and E. Culurciello, "Flattened convolutional neural networks for feedforward acceleration," *arXiv*, pp. 1–11, 2015.

[9] C. Fernández, R. Izquierdo, D. F. Llorca, and M. A. Sotelo, "Road curb and lanes detection for autonomous driving on urban scenarios," in *Proc. IEEE 17th Int*, pp. 1964–1969, Conf. Intell. Transp. Syst. (ITSC, October 2014.

[10] T. T. Duong, C. C. Pham, T. H. P. Tran, T. P. Nguyen, and J. W. Jeon, "Near real-time ego-lane detection in highway and urban streets," in *Proceedings of the IEEE International Conference on Consumer Electronics-Asia (ICCE-Asia 2016)*, (Korea, pp. 1–4.), pp. 26–28, Seoul, October 2016.

[11] A. Briot, P. Viswanath, and S. Yogamani, "Analysis of efficient cnn design techniques for semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 663–672, 2018.

[12] M. Yahiaoui, H. Rashed, L. Mariotti, G. Sistu, I. Clancy, L. Yahiaoui, V. R. Kumar, and S. Yogamani, "Fisheyemodnet: Moving object detection on surround-view cameras for autonomous driving," *arXiv preprint arXiv:1908.11789*, 2019.

[13] V. R. Kumar, S. Milz, C. Witt, M. Simon, K. Amende, J. Petzold, S. Yogamani, and T. Pech, "Near-field depth estimation using monocular fisheye camera: A semi-supervised learning approach using sparse lidar data," in *CVPR Workshop*, vol. 7, 2018.

[14] V. R. Kumar, S. Milz, C. Witt, M. Simon, K. Amende, J. Petzold, S. Yogamani, and T. Pech, "Monocular fisheye camera depth estimation using sparse lidar supervision," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pp. 2853–2858, IEEE, 2018.

[15] V. R. Kumar, S. A. Hiremath, M. Bach, S. Milz, C. Witt, C. Pinard, S. Yogamani, and P. Mäder, "Fisheyedistancenet: Self-supervised scale-aware distance estimation using monocular fisheye camera for autonomous driving," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 574–581, IEEE, 2020.

[16] V. R. Kumar, S. Yogamani, M. Bach, C. Witt, S. Milz, and P. Mader, "Unrectdepthnet: Self-supervised monocular depth estimation using a generic framework for handling common camera distortion models," *arXiv preprint arXiv:2007.06676*, 2020.

[17] N. Tripathi and S. Yogamani, "Trained trajectory based automated parking system using visual slam," *arXiv preprint arXiv:2001.02161*, 2020.

[18] M. Uřičář, P. Křížek, G. Sistu, and S. Yogamani, "Soilingnet: Soiling detection on automotive surround-view cameras," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 67–72, IEEE, 2019.

[19] M. Uricár, J. Ulicny, G. Sistu, H. Rashed, P. Krizek, D. Hurych, A. Vobecky, and S. Yogamani, "Desoiling dataset: Restoring soiled areas on automotive fisheye cameras," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 0–0, 2019.

[20] M. Uricar, G. Sistu, H. Rashed, A. Vobecky, V. R. Kumar, P. Krizek, F. Burger, and S. Yogamani, "Let's get dirty: Gan based data augmentation for camera lens soiling detection in autonomous driving," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 766–775, 2021.

[21] G. Sistu, I. Leang, S. Chennupati, S. Yogamani, C. Hughes, S. Milz, and S. Rawashdeh, "Neurall: Towards a unified visual perception model for automated driving," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 796–803, IEEE, 2019.

[22] S. Chennupati, G. Sistu, S. Yogamani, and S. A Rawashdeh, "Multinet++: Multi-stream feature aggregation and geometric loss strategy for multi-task learning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 0–0, 2019.

[23] M. Pöpperli, R. Gulagundi, S. Yogamani, and S. Milz, "Capsule neural network based height classification using low-cost automotive ultrasonic sensors," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, pp. 661–666, IEEE, 2019.

[24] H. Rashed, M. Ramzy, V. Vaquero, A. El Sallab, G. Sistu, and S. Yogamani, "Fusemodnet: Real-time camera and lidar based moving object detection for robust low-light autonomous driving," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pp. 0–0, 2019.

[25] K. El Madawi, H. Rashed, A. El Sallab, O. Nasr, H. Kamel, and S. Yogamani, "Rgb and lidar fusion based 3d semantic segmentation for autonomous driving," in *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pp. 7–12, IEEE, 2019.

[26] M. Uřičář., D. Hurych., P. Křížek, and S. Yogamani., "Challenges in

designing datasets and validation for autonomous driving," in *Proceedings of the 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications - Volume 5 VISAPP: VISAPP,*, pp. 653–659, INSTICC, SciTePress, 2019.

[27] S. Yogamani, C. Hughes, J. Horgan, G. Sistu, P. Varley, D. O'Dea, M. Uricar, S. Milz, M. Simon, K. Amende, C. Witt, H. Rashed, S. Chennupati, S. Nayak, S. Mansoor, X. Perrotton, and P. Perez, "Woodscape: A multi-task, multi-camera fisheye dataset for autonomous driving," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.

[28] G. J. Brostow, J. Fauqueur, and R. Cipolla, "Semantic object classes in video: A high- definition ground truth database," *Pattern Recognition Letters*, pp. 1–10, 2008.

[29] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez, "The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes," *CVPR*, pp. 3234–3243, 2016.

Author Biography

Ashok Dahal, Ph.D. *has been working as a Senior Computer Vision Software Engineer at Valeo North America Inc. (Valeo) since 2015. Dr. Dahal's work at Valeo involves the research and development of ADAS features which are sold to major automobile companies around the world. He is also involved in patent filing and Intellectual Property (IP) clearance process related to ADAS features and represents Valeo in various conferences and publications. Prior to joining Valeo, he completed his PhD from University of North Texas, Denton TX, with research focused on image and video analysis using computer vision and machine learning to improve colonoscopy videos capabilities.*

Eric Golab *is a Data Science Master's student at the University of Michigan, Ann Arbor MI. He previously attended Kettering University, in Flint MI, where he received a B.S. in Applied Physics and a B.S. in Applied Mathematics in 2018. Between 2015 and 2018 he worked for Valeo as a co-op, focusing on ADAS software development. This work became part of his undergraduate thesis. He is now computer vision specialist affiliated with Ford Motor Company.*

Rajender Garlapati *is Lane Sensing Feature Owner at Valeo North America Inc. Rajender is with Valeo since 2017 working as Lane sensing feature owner develops computer vision based system/function requirements, supports ADAS based feature architecture and evaluates ADAS feature performance with KPIs as well as in vehicles. He has over 12 years of experience in industrial embedded systems design and development and 5 years of experience in embedded software development. he has certification in embedded system design, development and Neural networks/deep learning.*

Varun Ravi Kumar *received a B.E. degree in 2015 and an M.Sc. degree in 2017 from TU Chemnitz, Germany. He is currently a Ph.D. student in Deep Learning for autonomous driving affiliated to TU Ilmenau and is currently working at Valeo. His research is mainly focused on the design of self-supervised perception algorithms using neural networks for self-driving cars. His expertise lies in depth and flow estimation for fisheye images and multi-task modeling. His focus also lies in semantic, motion segmentation, 2D and 3D object detection, and point cloud processing. He was awarded the* Deutschlandstipendium *for top-class international talent. He was also part of Udacity's first cohort of Self-Driving-Car Nanodegree in 2017.*

Senthil Yogamani *is an Artificial Intelligence architect and technical leader at Valeo. He leads the research and design of AI algorithms for various modules of autonomous driving systems. He has over 14 years of experience in computer vision and machine learning including 12 years of experience in industrial automotive systems. He is an author of over 90 publications and 60 patents with 1300+ citations. He serves in the editorial board of various leading IEEE automotive conferences including ITSC, IV and ICVES and advisory board of various industry consortia including Khronos, Cognitive Vehicles and IS Auto. He is a recipient of best associate editor award at ITSC 2015 and best paper award at ITST 2012.*

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

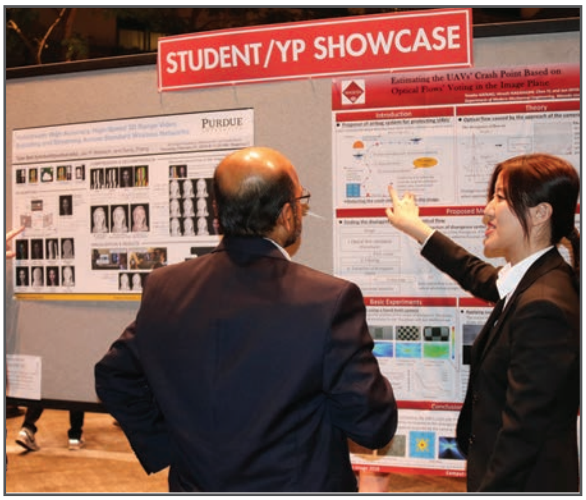

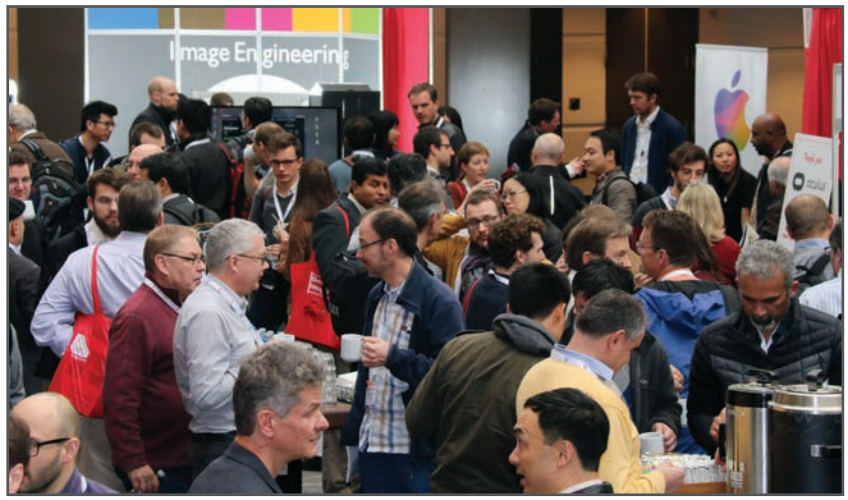

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org