

LiMoSeg: Real-time Bird's Eye View based LiDAR Motion Segmentation

Sambit Mohapatra¹, Mona Hodaei¹, Senthil Yogamani², Stefan Milz³, Heinrich Gotzig¹,
Martin Simon^{1,4}, Hazem Rashed¹ and Patrick Maeder⁴

¹Valeo, Germany

²Valeo, Ireland

³Spleenlab.ai, Germany

⁴TU Ilmenau, Germany

Keywords: Automated Driving, Point cloud processing, Motion Segmentation, Bird's Eye View Algorithms.

Abstract: Moving object detection and segmentation is an essential task in the Autonomous Driving pipeline. Detecting and isolating static and moving components of a vehicle's surroundings are particularly crucial in path planning and localization tasks. This paper proposes a novel real-time architecture for motion segmentation of Light Detection and Ranging (LiDAR) data. We use two successive scans of LiDAR data in 2D Bird's Eye View (BEV) representation to perform pixel-wise classification as static or moving. Furthermore, we propose a novel data augmentation technique to reduce the significant class imbalance between static and moving objects. We achieve this by artificially synthesizing moving objects by cutting and pasting static vehicles. We demonstrate a low latency of 8 ms on a commonly used automotive embedded platform, namely Nvidia Jetson Xavier. To the best of our knowledge, this is the first work directly performing motion segmentation in LiDAR BEV space. We provide quantitative results on the challenging SemanticKITTI dataset, and qualitative results are provided in <https://youtu.be/2aJ-cL8b0LI>.

1 INTRODUCTION

Autonomous Driving tasks such as perception which involves object detection (Rashed et al., 2021), (Dahal et al., 2021b), (Hazem et al., 2020), soiling detection (Uricar et al., 2021), (Das et al., 2020), road edge detection (Dahal et al., 2021a), weather classification (Dhananjaya et al., 2021), depth prediction (Varun et al., 2021b), (Ravi Kumar et al., 2018) is challenging due to the highly dynamic and interactive nature of surrounding objects in the automotive scenarios (Houben et al., 2021). Identification of the environmental objects as moving and static is crucial to achieving safe motion planning and navigation. An autonomous vehicles' route has to consider future coordinates and velocities of surrounding moving objects. In addition, this information is a critical source for simultaneous localization and mapping (SLAM) (Gallagher et al., 2021) and pose estimation (Kumar et al., 2020). As the vehicle is in motion, it is difficult to distinguish between background and other moving objects. Thus, motion segmentation requires estimation of the vehicle's ego-motion

and compensation to extract other moving objects in the scene. Motion cues can also be used to detect generic moving objects like animals which are difficult to train based on appearance due to their rare appearance in driving scenes and due to their diversity. Relative to appearance-based object detection and semantic segmentation, CNN-based motion segmentation approaches are relatively less mature (Ravi Kumar et al., 2021b), (Yahiaoui et al., 2019).

Autonomous vehicles are equipped with a variety of sensors to generate an understanding of environments. The most common ones are cameras and LiDAR. Although cameras provide rich color information, they face a lack of depth information and rely on illumination, making them vulnerable to poor illumination conditions such as nights or rainy days and are also prone to adversarial attacks (Sobh et al., 2021). However, providing accurate 3D depth (Varun et al., 2021a), (Varun et al., 2020) and geometric information (Ravi Kumar et al., 2021a) of the environment without dependency on weather and illumination is possible with LiDAR (Kumar et al., 2018). Considering the benefits of LiDAR data, we focus our ef-

forts towards motion segmentation in LiDAR point clouds, building upon the blocks presented in (Mohanapatra et al., 2021). A summary of the contributions of this work are listed below:

- We propose a novel method to implement real-time motion segmentation on LiDAR point clouds. First, we convert LiDAR 3D point clouds to 2D Bird's Eye View (BEV), then we classify each pixel of the BEV as static or motion. We demonstrate real-time inference on an embedded GPU platform.
- We introduce a novel residual computation layer that directly leverages the motion across frames to increase the disparity between pixel values for static and moving parts of the BEV frames.
- We introduce a data augmentation technique to simulate motion by selectively translating static objects across successive frames. The technique addresses the problem of significant class imbalance present in the dataset.

2 RELATED WORK

A variety of approaches have been proposed for moving semantic segmentation using vision (Barnes et al., 2018) (Patil et al., 2020) (McManus et al., 2013). Apart from vision-based methods, other approaches rely on the fusion of vision and LiDAR sensors (Rashed et al., 2019) (El Madawi et al., 2019) (Yan et al., 2014) (Postica et al., 2016). Using LiDAR sensors individually in order to perform semantic segmentation tasks has been taken into consideration recently (Cortinhal et al., 2020) (Li et al., 2020) (Milioto et al., 2019).

Motion segmentation can be performed by LiDAR-based methods based on clustering approaches such as (Dewan et al., 2016a) including point motion prediction by RANSAC and clustering objects. Vaquero et al. (Vaquero et al., 2017) performed motion segmentation after clustering vehicles points and matching objects in consecutive frames. Steinhauser et al. (Steinhauser et al., 2008) have devised another method to classify moving and non-moving objects using RANSAC and extract features from two sequential frames, although in some scenarios like when a vehicle is surrounded by many moving objects or dense trees, the RANSAC algorithm and feature extraction have failed.

In other studies, Wang et al. (Wang et al., 2012) have segmented objects that are able to move into different categories such as cars, bicycles, and pedestrians in laser scans of urban senses. Consistent tem-

poral information of consecutive LiDAR scans has been utilized with semantic classification and semantic segmentation approaches (Dewan and Burgard, 2020) (Dewan et al., 2017) which are developed based on motion vectors of rigid bodies that have been estimated by a flow approach on LiDAR scans (Dewan et al., 2016b). However, distinguishing scene flow from noise in the case of slowly moving objects can be a difficult task to perform. In urban scenarios, most semantic segmentation methods are able to recognize objects typically being in the pedestrians, bicyclists, and cars classes in point clouds (Alonso et al., 2020) (Milioto et al., 2019) (Wu et al., 2019) (Wu et al., 2018) (Biasutti et al., 2019) (Cortinhal et al., 2020). However, none of them distinguish between static and moving objects. There are some studies in order to distinguish moving objects. Yoon et al. (Yoon et al., 2019) propose a ray-tracing method including a clustering step to detect moving objects in LiDAR scans, which can, however, occasionally result in incomplete detection of objects or detection of static areas. Shi et al. (Shi et al., 2020) introduce a method based on utilizing sequential point clouds to achieve the prediction of moving objects. However, most of the used architectures were primarily developed for semantic segmentation and have a relatively large number of parameters.

3 PROPOSED METHOD

3.1 Input Data Prepration

Since a set of sequential point clouds is needed for motion segmentation, we group each frame in the SemanticKITTI (Behley et al., 2019) dataset with its past two frames (we use 2 past frames) in the sequence. The past frames are then motion compensated using the pose matrices provided in the dataset as described by the Eq. 1. The motivation behind using 2 past frames is to strike a balance between the accuracy of predictions and network size and latency. The number of encoders is directly proportional to the number of input frames.

$$F_{N-1}^N = P_N^{-1} \cdot (P_{N-1} \cdot F_{N-1}) \quad (1)$$

where

F_N : frame N

F_{N-1} : frame $N - 1$

P_N : Pose matrix of frame N

P_{N-1} : Pose matrix of frame $N - 1$

The set of motion-compensated frames are then converted to 2D BEV using the same process as de-




Figure 1: Input 3 frames superimposed - red pixels are present frame (Left), Ground truth motion mask (Middle), Predicted motion mask (Right).

scribed in (Mohapatra et al., 2021), and (Barrera et al., 2020). We limit the x and y ranges to (0, 48) and (-16, 16) meters, respectively, for BEV creation. Using a cell resolution (in BEV space) of 0.1 meters, we generate (480x320) sized BEV images for each LiDAR frame.

Unlike (Chen et al., 2021), we do not use the range image representation. Our reason is that range images are only better at short range due to their skewed aspect ratio. Neighboring pixels in range representations disregard the metric distance from the underlying points compared to BEV. Objects far away from the car are barely visible in a range image. Furthermore, range images are affected by even partial occlusion. BEV representation overcomes these problems to some extent (particularly for semi-occluded objects) while presenting the benefits of 2D representation. Another key advantage of BEV representation is that reconstruction of 3D points is a simple matter of looking up the row and column indices of the pixels and multiplying by the cell resolution. Furthermore, most downstream applications such as motion planning are made on a grid-based BEV space, and hence predictions available directly in BEV space reduce the number of interconversions.

3.2 Data Augmentation

Due to the rather difficult and expensive process of collecting and annotating LiDAR data, data augmentation has been a key technique used to increase the training set size and also allow better generalization of the network to different scenarios. One of the most commonly used techniques is sampling-based augmentation or ground truth augmentation as described

in Second algorithm (Yan et al., 2018). The idea is to copy objects from some frames and paste them into others, increasing the number and type of objects. We use this idea but modify it to fit our case of generating artificially moving objects in frames.

For each frame (3D LiDAR) with no moving objects (no points marked as moving), we collect all the points belonging to class cars. A uniform random value then translates these points along x and y axes for N successive frames, and the translations increase along the x-axis in each frame to produce a notion of motion. Experimentally, we found N=4 to perform best. The transformed points are then concatenated to the rest of the points, and their labels are marked as moving cars. We do not apply this technique to frames with motion objects to avoid clutter due to overlap between the synthetic objects and the actual objects in the frame. Though simple, this method enables us to make better use of a sizeable portion of the dataset that has no or very few numbers of LiDAR points marked as moving objects.

3.3 Network Architecture

Our goal is design a a pixel-wise prediction model which operates in BEV space and it is also efficient with very low latency. BEVDetNet (Mohapatra et al., 2021) is a recent efficient model which operates on BEV space. It produces outputs in the same spatial resolution as input and has a head that does binary keypoint classification. However, for predicting relative motion between frames and classifying each pixel as moving or static, we adapt this architecture. We inherit the building blocks from BEVDetNet (Mohapatra et al., 2021) and build a multi-encoder joint-

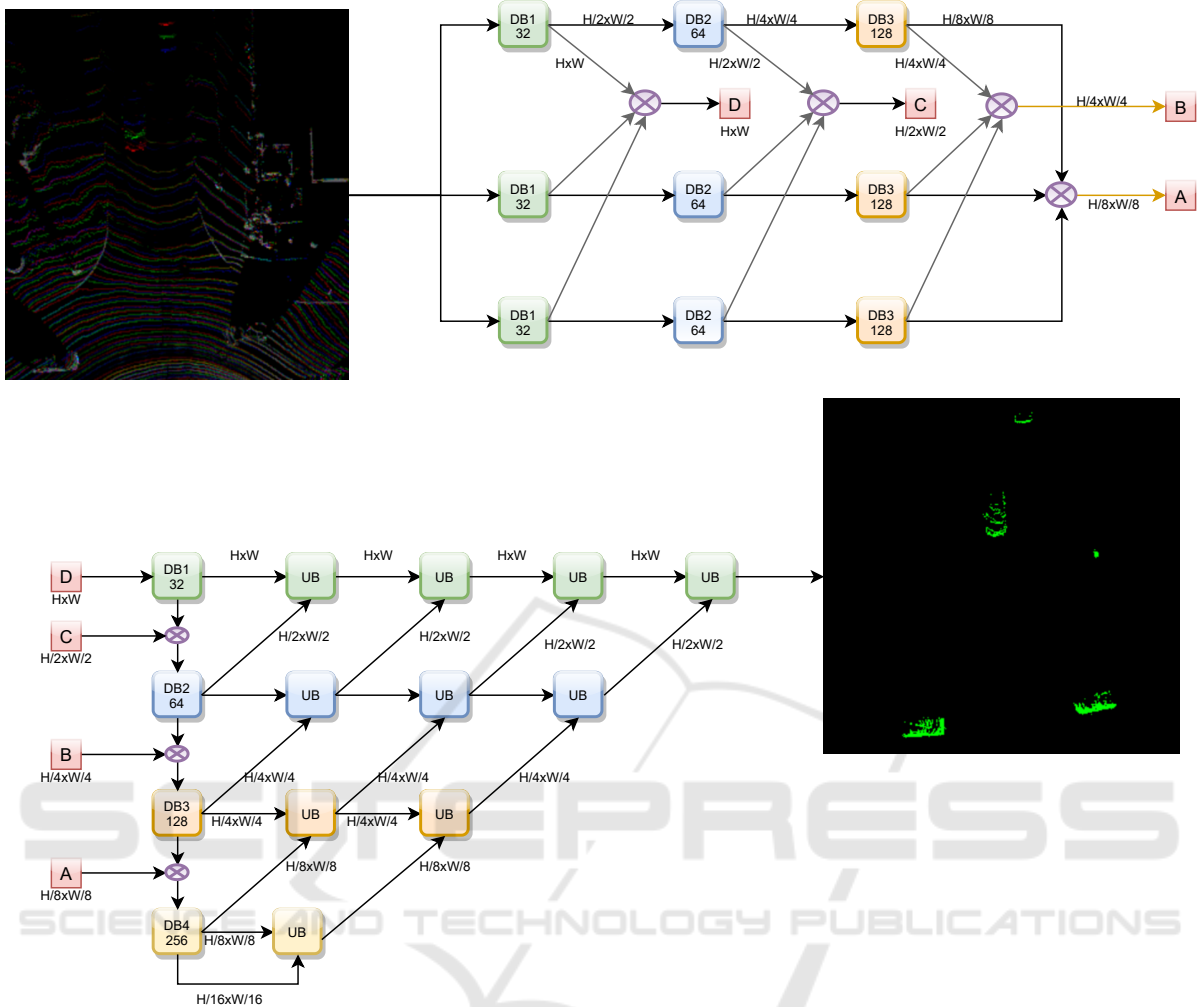


Figure 2: Overall architecture of LiMoSeg with individual encoders (top) and joint encoder and decoder parts (bottom). The three input BEV frames are superimposed as red-green-blue channels. The red channel is the present frame and the green and blue channels are past two frames in sequence.

decoder architecture as shown in Figure 2.

The feature extraction blocks are called Down-sampling Blocks (DB), as seen in Figure 3. They use 5×5 and 3×3 convolutions to capture features at different scales and also successively reduce the spatial resolution of input using a max-pooling operation at the end. The Upsampling Blocks (UB), as seen in Figure 3 are used to increase the spatial resolution of inputs and serve to produce final output at the same spatial dimension as the input. It consists of a transposed convolution followed by a single convolutional block. We use the ReLU activation function throughout the network. Since we have 3 input BEV images (present frame and past 2 frames), we have 3 individual encoders consisting of 3 DB blocks. The individual encoders compute per-input features. Features from multiple stages of the network are then

collected for all three encoders and then fused using a concatenation and multiplication-based fusion approach. A joint feature computation chain consisting of 4 DB blocks then computes joint features upon the pre-computed individual features from each encoder. The idea is, individual encoders compute features for objects while the joint encoder computes features that capture the interaction between objects from all three streams. These are essentially the features that capture the relative displacement between objects due to motion. While concatenation is used as a most common feature fusion approach, we augment it by explicitly multiplying features channel-wise from corresponding stages to compute a set of features that forms a loose correlation between similar features across the three channels. We then concatenate these features to the rest of the concatenated features. We




Figure 3: Downsampling Block (DB) (Left) and Upsampling Block (UB) (Right).

use only 3 DB blocks in the individual encoders since their primary job is to compute the low-level features that characterize objects. The joint encoder has 4 DB blocks to allow computation of sufficiently complex features to capture motion. Our motivation was to limit the number of parameters as much as possible.

3.3.1 Residual Computation Layer

Residual layers have proven to improve motion segmentation significantly as demonstrated by (Chen et al., 2021). They compute the difference between successive motion-compensated frames and then normalize it. This produces a disparity map between the moving and static parts of the two frames. However, we multiply motion-compensated successive frames to generate residuals. Since static objects will have overlaps (some of the parts at least) across successive frames, the residual gets large values in such parts. Moving objects occupy different spatial locations across successive frames (with some overlap depending on the amount of motion). Therefore, such locations become 0 due to the multiplication as seen in Figure 4. This creates a much more significant disparity in static and moving parts of the frame and provides a weak attention mechanism to the network. We normalize again after residual computation.

4 RESULTS

We follow the standard training and validation split for the semantic KITTI dataset and keep sequence 8 for validation. All other sequences between [00-10] are converted into BEV, and the motion annotations are also converted into BEV. It was experimentally found that due to the significant class imbalance between frames with motion points and frames without motion points, using all the frames affects the network’s performance negatively by acting as a bias towards classifying points as static more often. There-




Figure 4: Three successive frames show a moving object at the center in red-green-blue (Left), Computed residuals with static parts in bright, and mobile objects region in dark at the center (Right).

Table 1: Comparison of accuracy and inference latency per frame in BEV for motion segmentation.

<i>Method</i>	<i>Inference Latency (ms)</i>	<i>IoU (moving class)</i>
RangeNet++ (Chen et al., 2021)	45	39.5
MINet (Chen et al., 2021)	24	36.9
SalsaNext (Chen et al., 2021)	41	53.4
LiMoSeg (Ours)	8	52.6

fore, we only convert those frames to BEV, which have at least 20 motion points in them. We train with a batch size of 12 for only 30 epochs. Since this is a segmentation task, we use weighted cross-entropy loss as our loss function as described by equation 2, where y_c is the ground truth class and \hat{y}_c is the predicted class for each pixel in the BEV.

$$L_{motion} = - \sum_{c=1}^C w_c y_c \log \hat{y}_c, \quad w_c = \frac{1}{\log(f_c + \epsilon)} \quad (2)$$

Table 2: Ablation study of different architectural and residual computation settings.

<i>Architecture</i>	<i>ch. wise mul</i>	<i>ch. wise sub</i>	IoU
Single encoder	✓	✗	43.20
Single encoder	✗	✓	38.36
Single encoder with semantics	✓	✗	39.23
Multiple encoders	✓	✗	43.45
Multiple encoders with joint features	✓	✗	52.60

Table 3: Ablation study of inference models illustrating the precision, accuracy and size.

<i>Precision</i>	<i>IoU</i>	<i>Latency (ms)</i>	<i>Size (MB)</i>
FP32	52.60	8	35.0
FP16	51.40	3	15.3
INT8	48.07	2	8.0

For evaluating the performance, we use the intersection-over-union (Everingham et al., 2010) metric as is commonly used by similar methods such as (Chen et al., 2021). The evaluation code is taken directly from SemanticKITTI. We evaluate our model in the BEV space since most of the algorithms downstream from the perception task, such as path planning, are carried out in the BEV space. We report Intersection over Union (IoU) for moving class.

To prove the real-time capabilities of our proposed architecture, we run the inference on an Nvidia Jetson Xavier AGX development kit, a commonly used automotive embedded systems platform for deep learning applications. As is shown in Table 1, we achieve an impressive inference latency of 8ms (inference speed of 125 Frames Per Second (FPS)). In terms of accuracy, we are slightly behind SalsaNext (Cortinhal et al., 2020). We haven’t performed extensive hyperparameter tuning or data augmentation.

4.1 Ablation Study

We perform extensive ablation studies using different modifications to our architecture and input data representation. As can be seen in Table 2, we evaluate using both a single encoder-decoder network as well as with multiple encoders and joint decoder architecture, which is more common in optical flow-based approaches such as (Ilg et al., 2017). However, surprisingly enough, the single encoder-decoder architecture achieves better accuracy than a multi-encoder approach where joint features were not computed. We explain that each encoder during training learns

features local to its input BEV frame for the multi-encoder approach. However, due to the absence of any DB blocks for joint feature learning, sufficient joint features are not learned, which leads to a lower representation of motion in feature space.

We also experiment with adding full semantic segmentation masks to individual input BEVs. This, however, does not seem to offer any benefits. The network seems to be biased towards classifying all instances of objects as positive for motion in this case. This could be directly attributed to the full semantic masks that do not differentiate between mobile and static objects. Furthermore, we evaluate another commonly used residual computation approach like subtraction and find that the proposed approach to multiply performs best.

We experiment with reduced precision inference as shown in Table 3 at 16-bit floating point (FP16) and 8-bit integer (INT8) and demonstrate that performance is not degraded by a large amount even at reduced precision. This is directly relevant for real-time embedded system based deployment.

5 CONCLUSION

In this paper, we demonstrated an algorithm that incrementally builds upon an existing network for object detection to do motion segmentation so that it can be added as an additional task in a multi-task learning framework. We proposed a residual computation layer that exploits the disparity between the static and mobile parts of two successive motion-compensated frames. We also proposed a data augmentation technique that greatly improves the class imbalance between static and mobile points present in the SemanticKITTI dataset. We observed that motion segmentation in BEV space is not a straightforward task due to sparsity of the 3D points and a relatively small cross-section of several traffic elements like pedestrians and bicyclists.

ACKNOWLEDGEMENT



We are funded by the Electronic Components and Systems for European Leadership Joint Undertaking grant No 826655 receiving support from the European Union’s Horizon 2020 research and innovation programme. Further partial funding is provided by the German Federal Ministry of Education and Research.

REFERENCES

- Alonso, I., Riazuelo, L., Montesano, L., and Murillo, A. C. (2020). 3D-MiniNet: Learning a 2D Representation from Point Clouds for Fast and Efficient 3D LIDAR Semantic Segmentation. *IEEE Robotics and Automation Letters*, 5(4):5432–5439.
- Barnes, D., Maddern, W., Pascoe, G., and Posner, I. (2018). Driven to Distraction: Self-Supervised Distractor Learning for Robust Monocular Visual Odometry in Urban Environments. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1894–1900. IEEE.
- Barrera, A., Guindel, C., Beltrán, J., and García, F. (2020). BirdNet+: End-to-End 3D Object Detection in LiDAR Bird’s Eye View. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE.
- Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., and Gall, J. (2019). SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*.
- Biasutti, P., Lepetit, V., Aujol, J.-F., Brédif, M., and Bugeau, A. (2019). LU-Net: An Efficient Network for 3D LiDAR Point Cloud Semantic Segmentation Based on End-to-End-Learned 3D Features and U-Net. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0.
- Chen, X., Li, S., Mersch, B., Wiesmann, L., Gall, J., Behley, J., and Stachniss, C. (2021). Moving Object Segmentation in 3D LiDAR Data: A Learning-Based Approach Exploiting Sequential Data. *arXiv preprint arXiv:2105.08971*.
- Cortinhal, T., Tzelepis, G., and Aksoy, E. E. (2020). SalsaNext: Fast, Uncertainty-aware Semantic Segmentation of LiDAR Point Clouds for Autonomous Driving. *arXiv preprint arXiv:2003.03653*.
- Dahal, A., Golab, E., Garlapati, R., Ravi Kumar, V., and Yogamani, S. (2021a). RoadEdgeNet: Road Edge Detection System Using Surround View Camera Images. In *Electronic Imaging*.
- Dahal, A., Kumar, V. R., Yogamani, S., and Eising, C. (2021b). An online learning system for wireless charging alignment using surround-view fisheye cameras. *arXiv preprint arXiv:2105.12763*.
- Das, A., Křížek, P., Sistu, G., Bürger, F., Madasamy, S., Uříčář, M., Ravi Kumar, V., and Yogamani, S. (2020). TiledSoilingNet: Tile-level Soiling Detection on Automotive Surround-view Cameras Using Coverage Metric. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6. IEEE.
- Dewan, A. and Burgard, W. (2020). DeepTemporalSeg: Temporally Consistent Semantic Segmentation of 3D LiDAR Scans. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2624–2630. IEEE.
- Dewan, A., Caselitz, T., Tipaldi, G. D., and Burgard, W. (2016a). Motion-based detection and tracking in 3D LiDAR scans. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 4508–4513. IEEE.
- Dewan, A., Caselitz, T., Tipaldi, G. D., and Burgard, W. (2016b). Rigid scene flow for 3D LiDAR scans. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1765–1770. IEEE.
- Dewan, A., Oliveira, G. L., and Burgard, W. (2017). Deep Semantic Classification for 3D LiDAR Data. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3544–3549. IEEE.
- Dhananjaya, M. M., Ravi Kumar, V., and Yogamani, S. (2021). Weather and light level classification for autonomous driving: Dataset, baseline and active learning. *arXiv preprint arXiv:2104.14042*.
- El Madawi, K., Rashed, H., El Sallab, A., Nasr, O., Kamel, H., and Yogamani, S. (2019). Rgb and lidar fusion based 3d semantic segmentation for autonomous driving. In *2019 IEEE Intelligent Transportation Systems Conference (ITSC)*, pages 7–12. IEEE.
- Everingham, M., Van Gool, L., Williams, C. K., Winn, J., and Zisserman, A. (2010). The PASCAL Visual Object Classes (VOC) Challenge. *International journal of computer vision*, 88(2):303–338.
- Gallagher, L., Kumar, V. R., Yogamani, S., and McDonald, J. B. (2021). A hybrid sparse-dense monocular slam system for autonomous driving. In *Proc. of ECMR*, pages 1–8. IEEE.
- Hazem, R., Mohamed, E., Sistu, Ganesh and, V. R. K., Eising, C., El-Sallab, A., and Yogamani, S. (2020). FisheyeYOLO: Object Detection on Fisheye Cameras for Autonomous Driving. *Machine Learning for Autonomous Driving NeurIPS 2020 Virtual Workshop*.
- Houben, S., Abrecht, S., Akila, M., Bär, A., et al. (2021). Inspect, Understand, Overcome: A Survey of Practical Methods for AI Safety. *CoRR*, abs/2104.14235.
- Ilg, E., Mayer, N., Saikia, T., Keuper, M., Dosovitskiy, A., and Brox, T. (2017). FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2462–2470.
- Kumar, V. R., Hiremath, S. A., Bach, M., Milz, S., Witt, C., Pinard, C., Yogamani, S., and Mäder, P. (2020). FisheyeDistancenet: Self-supervised scale-aware distance estimation using monocular fisheye camera for autonomous driving. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 574–581.
- Kumar, V. R., Milz, S., Witt, C., Simon, M., Amende, K., Petzold, J., Yogamani, S., and Pech, T. (2018). Monocular fisheye camera depth estimation using sparse lidar supervision. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2853–2858.
- Li, S., Chen, X., Liu, Y., Dai, D., Stachniss, C., and Gall, J. (2020). Multi-scale Interaction for Real-time LiDAR Data Segmentation on an Embedded Platform. *arXiv preprint arXiv:2008.09162*.
- McManus, C., Churchill, W., Napier, A., Davis, B., and Newman, P. (2013). Distraction suppression for

- vision-based pose estimation at city scales. In *2013 IEEE international conference on robotics and automation*, pages 3762–3769. IEEE.
- Milioto, A., Vizzo, I., Behley, J., and Stachniss, C. (2019). RangeNet ++: Fast and Accurate LiDAR Semantic Segmentation. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4213–4220. IEEE.
- Mohapatra, S., Yogamani, S., Gotzig, H., Milz, S., and Mader, P. (2021). BEVDetNet: Bird's Eye View LiDAR Point Cloud based Real-time 3D Object Detection for Autonomous Driving. *arXiv preprint arXiv:2104.10780*.
- Patil, P. W., Biradar, K. M., Dudhane, A., and Murala, S. (2020). An End-to-End Edge Aggregation Network for Moving Object Segmentation. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8149–8158.
- Postica, G., Romanoni, A., and Matteucci, M. (2016). Robust Moving Objects Detection in Lidar Data Exploiting Visual Cues. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1093–1098. IEEE.
- Rashed, H., Mohamed, E., Sistu, G., Ravi Kumar, V., Eising, C., El-Sallab, A., and Yogamani, S. (2021). Generalized Object Detection on Fisheye Cameras for Autonomous Driving: Dataset, Representations and Baseline. In *Proceedings of the Winter Conference on Applications of Computer Vision*, pages 2272–2280.
- Rashed, H., Ramzy, M., Vaquero, V., El Sallab, A., Sistu, G., and Yogamani, S. (2019). FuseMODNet: Real-Time Camera and LiDAR based Moving Object Detection for robust low-light Autonomous Driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0.
- Ravi Kumar, V., Klingner, M., Yogamani, S., Milz, S., Fingscheidt, T., and Mader, P. (2021a). Syndistnet: Self-supervised monocular fisheye camera distance estimation synergized with semantic segmentation for autonomous driving. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 61–71.
- Ravi Kumar, V., Milz, S., Witt, C., and Yogamani, S. (2018). Near-field depth estimation using monocular fisheye camera: A semi-supervised learning approach using sparse LiDAR data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*, volume 7.
- Ravi Kumar, V., Yogamani, S., Rashed, H., Sitsu, G., Witt, C., Leang, I., Milz, S., and Mäder, P. (2021b). Omnidet: Surround view cameras based multi-task visual perception network for autonomous driving. *IEEE Robotics and Automation Letters*, 6(2):2830–2837.
- Shi, H., Lin, G., Wang, H., Hung, T.-Y., and Wang, Z. (2020). SpSequenceNet: Semantic Segmentation Network on 4D Point Clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4574–4583.
- Sobh, I., Hamed, A., Kumar, V. R., and Yogamani, S. (2021). Adversarial attacks on multi-task visual perception for autonomous driving. *arXiv preprint arXiv:2107.07449*.
- Steinhaus, D., Ruepp, O., and Burschka, D. (2008). Motion segmentation and scene classification from 3D LiDAR data. In *2008 IEEE intelligent vehicles symposium*, pages 398–403. IEEE.
- Uricar, M., Sistu, G., Rashed, H., Vobecky, A., Ravi Kumar, V., Krizek, P., Burger, F., and Yogamani, S. (2021). 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*, pages 766–775.
- Vaquero, V., Del Pino, I., Moreno-Noguer, F., Sola, J., Sanfeliu, A., and Andrade-Cetto, J. (2017). Deconvolutional networks for point-cloud vehicle detection and tracking in driving scenarios. In *2017 European Conference on Mobile Robots (ECMR)*, pages 1–7. IEEE.
- Varun, R. K., Klingner, M., Yogamani, S., Bach, M., Milz, S., Fingscheidt, T., and Mäder, P. (2021a). SVDistNet: Self-supervised near-field distance estimation on surround view fisheye cameras. *IEEE Transactions on Intelligent Transportation Systems*.
- Varun, R. K., Yogamani, S., Bach, M., Witt, C., Milz, S., and Mäder, P. (2020). UnRectDepthNet: Self-Supervised Monocular Depth Estimation using a Generic Framework for Handling Common Camera Distortion Models. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*.
- Varun, R. K., Yogamani, S., Milz, S., and Mäder, P. (2021b). FisheyeDistanceNet++: Self-Supervised Fisheye Distance Estimation with Self-Attention, Robust Loss Function and Camera View Generalization. In *Electronic Imaging*.
- Wang, D. Z., Posner, I., and Newman, P. (2012). What could move? Finding cars, pedestrians and bicyclists in 3D laser data. In *2012 IEEE International Conference on Robotics and Automation*, pages 4038–4044. IEEE.
- Wu, B., Wan, A., Yue, X., and Keutzer, K. (2018). SqueezeSeg: Convolutional Neural Nets with Recurrent CRF for Real-Time Road-Object Segmentation from 3D LiDAR Point Cloud. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE.
- Wu, B., Zhou, X., Zhao, S., Yue, X., and Keutzer, K. (2019). SqueezeSegV2: Improved Model Structure and Unsupervised Domain Adaptation for Road-Object Segmentation from a LiDAR Point Cloud. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 4376–4382. IEEE.
- Yahiaoui, M., Rashed, H., Mariotti, L., Sistu, G., Clancy, I., Yahiaoui, L., Ravi Kumar, V., and Yogamani, S. (2019). FisheyeModNet: Moving object detection on Surround-View Cameras for Autonomous Driving. *arXiv preprint arXiv:1908.11789*.
- Yan, J., Chen, D., Myeong, H., Shiratori, T., and Ma, Y. (2014). Automatic Extraction of Moving Objects from Image and LIDAR Sequences. In *2014 2nd International Conference on 3D Vision*, volume 1. IEEE.
- Yan, Y., Mao, Y., and Li, B. (2018). Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337.
- Yoon, D., Tang, T., and Barfoot, T. (2019). Mapless Online Detection of Dynamic Objects in 3D Lidar. In *2019 16th Conference on Computer and Robot Vision (CRV)*, pages 113–120. IEEE.