LRPD: Long Range 3D Pedestrian Detection Leveraging Specific Strengths of LiDAR and RGB

Michael Fürst[1] and Oliver Wasenmüller[1] and Didier Stricker[1]

[1]DFKI - German Research Center for Artificial Intelligenc

`firstname.lastname@dfki.de`

Abstract— While short range 3D pedestrian detection is sufficient for emergency breaking, long range detections are required for smooth breaking and gaining trust in autonomous vehicles. The current state-of-the-art on the KITTI benchmark performs suboptimal in detecting the position of pedestrians at long range. Thus, we propose an approach specifically targeting long range 3D pedestrian detection (LRPD), leveraging the density of RGB and the precision of LiDAR. Therefore, for proposals, RGB instance segmentation and LiDAR point based proposal generation are combined, followed by a second stage using both sensor modalities symmetrically. This leads to a significant improvement in mAP on long range compared to the current state-of-the art. The evaluation of our LRPD approach was done on the pedestrians from the KITTI benchmark.

I. INTRODUCTION

Vulnerable Road User (VRU) detection is a major challenge to be solved for autonomous vehicles. The vehicle should avoid all obstacles including pedestrians. To successfully achieve this, a dependable 3D detection is required. Despite visually appealing results, current state-of-the-art approaches show suboptimal performance when numerically evaluated. One key challenge is the detection of pedestrians at long ranges. Between 20 to 50 meters is a typical breaking distance in cities at $0.3g$ breaking acceleration, leading to a comfortable driving experience. Beyond comfort, long range detection also gives an edge on security, since it increases the amount of time available to decision making and temporal fusion.

Even though long range 3D pedestrian detection is of significant importance for autonomous vehicles and automated driving functions, there has been little research done on that topic yet. We are the first to have an in-depth look at the performance of existing state-of-the-art approaches at long range. Comparing their mAP at long ranges, we found strategies to improve it, by using precise RGB segmentation in combination with LiDAR distance information. Other approaches do not have this strong focus on RGB in their proposals, even though RGB contains more information on far away objects. Our analysis goes beyond the official KITTI benchmark [1] and we show that the approaches have different performance characteristics, while achieving the same KITTI score.

Since all existing approaches have deficits in long range 3D pedestrian detection, we propose:

- A new proposal generation stage, leveraging density of RGB images and precision of LiDAR data, followed by
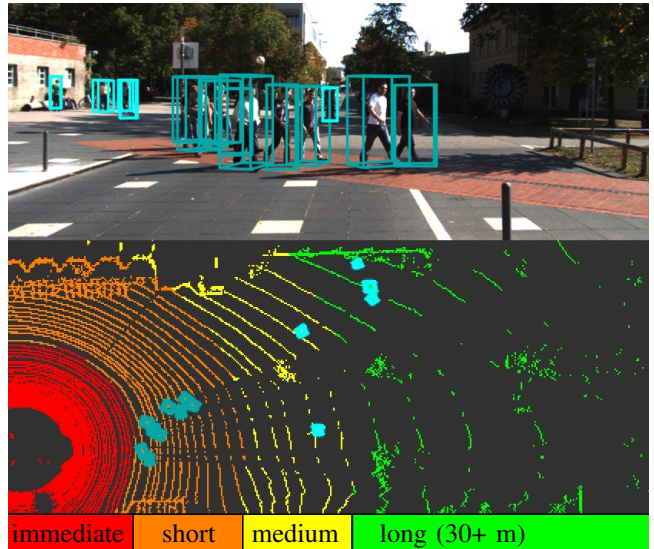

Fig. 1. Immediate (0-10m), short (10-20m), medium (20-30m) and long range ($\geq$ 30m) can be used to analyse a detectors performance. Recognizing its importance, we focus on long range and outperform the current state-of-the-art.

a symmetrical refinement stage,
- a distance based proposal augmentation schema to further enhance the performance of algorithms,
- an evaluation of existing approaches and our LRPD approach with a significant improvement in mAP over the current state-of-the-art on long range pedestrian detection.

II. RELATED WORK

3D detection of objects is a key technology required for autonomous driving and automated driving functions. Recently, there have been significant advances in this field.

PointPillars [2] uses only LiDAR information in an SSD [3] like network architecture. One key contribution of PointPillars is the introduction of pillars instead of voxels as in VoxelNet [4]. A feature embedding is computed for every non empty pillar, which is then used in a 2D BEV feature map, enabling efficient 2D convolutions. The 2D convolution backbone is an FPN [5] which is then used for an SSD style decoder head. We consider PointPillars as LiDAR only single-stage 3D detector.

Most approaches however are multi stage approaches, using an ROI-pooling approach derived from the R-CNN

series [6], [7], [8]. F-PointNet [9] leverages the great performance of 2D detectors, by predicting 2D bounding boxes, which are used to define a frustum in 3D space. All points from the LiDAR, which fall inside the frustum, are then used in a PointNet [10] which predicts an instance segmentation of the points. Only points which are segmented as foreground are then used for the final prediction of the bounding box by a PointNet. F-PointNet relies on 2D bounding box prediction in the RGB image to reduce the search space for objects and then does the final prediction of objects based only on the frustum LiDAR pointcloud using PointNets.

Whereas F-PointNet [9] uses classical 2D proposal based ROI pooling, AVOD [11] uses 3D ROI pooling. The architecture is similar to a Faster-R-CNN [8], but uses a static 3D anchor grid to crop and resize the features before fusing them and applying a fully connected network as an RPN on them. The outputs of this RPN are then used to crop and resize the features again, fuse them and finally estimate the bounding box by using a fully connected layer. AVOD is a two-stage 3D detector using RGB and LiDAR symmetrically in the RPN and refinement stage.

In contrast to previous approaches, IPOD [12] uses an unconventional proposal strategy focused on semantic segmentation. RGB segmentation is used to find the relevant points in the pointcloud and samples boxes around them. The proposal stage has influenced our proposal stage, since it can yield proposals from only one inlier point. However, IPOD only uses LiDAR pointclouds for its refinement stage.

State-of-the-Art is typically compared on the KITTI Benchmark [1]. It consists out of 7481 training images and 7518 testing images. Containing 4487 pedestrian instances in 1779 images for the training set. The top performing approaches are shown in Table 1. For our in-depth analysis we use AVOD [11], PointPillars [2] and F-PointNet [9], since they are the best performing open source algorithms and represent the common algorithm categories. Those approaches can be differentiated into single or two stage approaches, LiDAR only or RGB + LiDAR and the way how they leverage 2D information and 2D detection performance.

We analyze the performance by measuring the mAP for different ranges. The mAP is computed by limiting the range, in which detections are considered, to an interval. As can be seen in Figure 2, the performance of approaches which have a focus on LiDAR for the final predictions (F-PointNet and PointPillars) have a significant drop in performance for long range when compared to approaches like AVOD which use both sensors. VMVS [13], STD [14], IPOD [12], F-ConvNet [15] and TANet [16] were not evaluated in detail, due to the lack of source code available or similarity to other approaches. VMVS, F-ConvNet and TANet are conceptually similar to AVOD, F-PointNet or PointPillars.

By analysing the KITTI dataset [17], we found a few indications which support our observations. On one hand, the number of LiDAR points for objects at long range is very low. On average an object further away than 30 meters has only 18 Points. Challenging objects consist of 5.5 points only, which makes distinction between noise, poles and pedestrians almost impossible. On the other hand, in
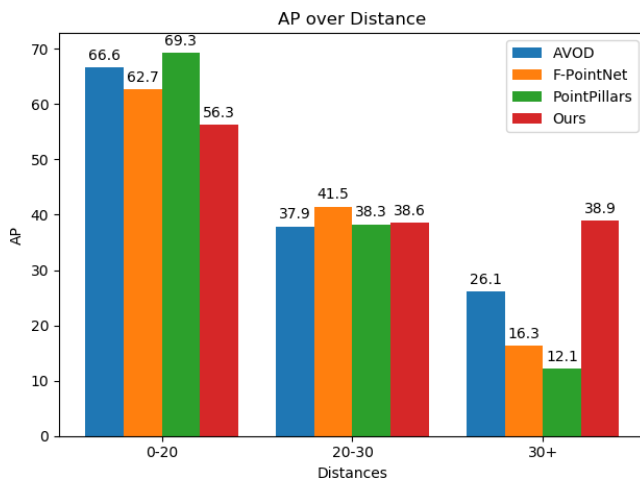

Fig. 2. Mean Average Precision (mAP) on KITTI BEV *easy* for pedestrians in 3D plotted for different ranges, using available state-of-the-art and our LRPD (with grounding) approach. We outperform all approaches by a margin on long range (30+ m), while maintaining competitive on short and medium range.

the RGB image the same objects consist of on average 706 pixels, which is sufficient for a classification in most cases. The differentiation between the pedestrian and any other pole like object can only be made by using RGB as a secondary input. PointPillars using only LiDAR is at a clear disadvantage. Similarly, F-PointNet relying on pointclouds for the refinement stage. AVOD achieves the best results, what can be attributed to the usage of RGB and LiDAR throughout all stages of the network. In other domains like scene flow estimation the potential of LiDAR and RGB fusion was already shown. [18] used LiDAR to improve the matching, interpolation and consistency check by constraining each step of their pipeline with LiDAR measurements. To the best of our knowledge there exists no algorithm, which specifically targets 3D pedestrian detection at long range. Thus, we present for the first time an approach targeting long range pedestrian detection from sparse LiDAR points and RGB data.

III. METHODOLOGY

From our analysis we identify the low number of LiDAR points per object as a problem for long range detection. Our approach, called LRPD, solves this problem. We have discovered, that a strong focus on RGB in the proposal stage and symmetrical fusion of LiDAR and RGB only in the second stage achieves best performance on long range detection tasks. All approaches presented in Section II, are either symmetrical in both stages or do not use both LiDAR and RGB in their second stage.

Based on an instance segmentation, our approach generates a proposal. The proposal is used for ROI cropping of the image and the pointcloud data, before feeding them through an encoder. Finally, the two resulting feature vectors are concatenated and a fully connected layer computes the objectness and refinements (Figure 3).
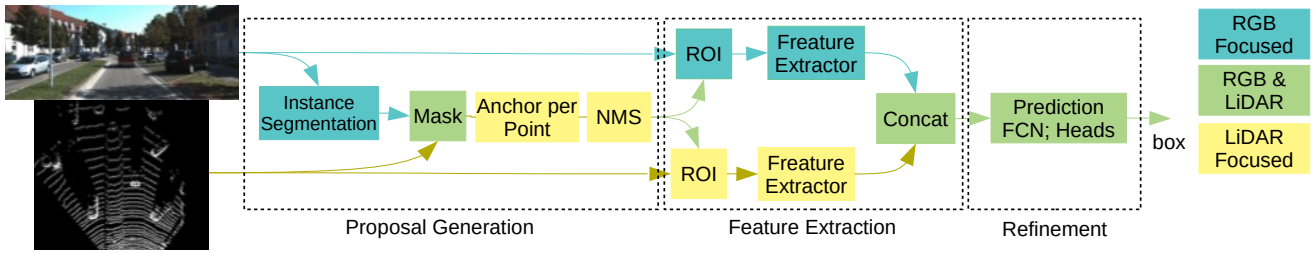
Fig. 3. The RGB image and LiDAR pointcloud are used for a proposal generation, then ROI cropped. The crops are used in a feature encoder, concatenated and finally the refinements for the proposal are predicted. The proposal is generated by applying Mask-RCNN instance masks to the LiDAR, creating a box around every foreground point and then using non-maximum suppression (NMS) to reduce the number of proposals.

A. Proposal Generation

Our proposal generation is a four step process inspired by IPOD [12], starting with segmentation, then masking the LiDAR with those instances, assigning a box to each positive point and finally applying non-maximum suppression (NMS). Even though our proposal pipeline looks similar to IPOD it varies in many details.

In contrast to IPOD – which uses semantic segmentation – we use an instance segmentation network, e.g. Mask-RCNN [19], which has shown strong results on KITTI and many other datasets. Instance segmentation should be preferred over semantic segmentation for two reasons. The latter often miss-classifies the area in-between two pedestrians, separated by a few pixels, as pedestrian. This leads to an unnecessary large amount of positive points resulting in a lot of unnecessary proposals. The other advantage of instance segmentation is that non-maximum suppression can be limited per instance, resulting in at least two proposals for two instances. With semantic segmentation this cannot be ensured.

After computing the instance masks, we apply them to the pointcloud to obtain the points related to an instance. This is done by projecting the pointcloud into the image space and assigning the points to the instance mask in which they fall or discard them if they are not inside any instance mask. The instance information is then applied back onto the original 3D pointcloud.

Now all 3D points that have an instance mask assigned are converted into a 3D bounding box proposal. The proposal has a width, height and length, defined by the average box on the dataset, and the x, y, z centeroid, defined by the point, used to generate it and the instance mask id is also stored.

Finally, all proposals undergo non-maximum suppression with a birds-eye-view (BEV) intersection over union (IoU) using a threshold of 0.5. However, our non-maximum suppression (NMS) harvests the knowledge about instances intelligently. The score for the NMS is defined by the number of inlier points with the same instance mask id a box has. This reduces the number of proposals greatly to 1-5 proposals per instance detected. With these changes in the proposal generation, the step of ambiguity reduction is not necessary, since we get a small amount of proposals per instance.

If the instance segmentation detects an object and the object contains at least one point, a proposal at the correct position is generated. Therefore, the main limiting factor of the proposal generation stage is the performance of the instance segmentation network.

B. Proposal Augmentation for Training

Besides the ability to generate proposals reliably from as little as one LiDAR point, proposal augmentation is a factor for the performance of our approach. Since the proposal stage often generates only one correct proposal per instance, there is little variance in proposal offsets. For the best performance of our refinement network, we increase variance by exploring various approaches for proposal augmentation. Random 3D displacement of the boxes and *grounding* has the most impact on performance. Random displacement directly increases variance in the proposals, whereas grounding simplifies the refinement task itself, reducing the necessity for large variance in the training data.

As simple data augmentation we use random displacement. All proposals are randomly displaced multiple times and a single copy of the original proposal is kept. After displacement a check is done, if the proposal is still close enough to the ground truth. This introduces what we call *close negative proposals*. Close negative proposals are close to a prediction, but not within the threshold.

At long range the proposals have less points. Therefore, using point based sampling, they might be significantly too high or low. To alleviate this issue we introduced what we call *grounding*. Grounding generates a pillar around the proposal. In this context pillar we find the lowest point and assume it to be the ground. The ground is preferred over the top, since poles stick out beyond the head of a pedestrian.

However, as objects get closer to the vehicle or are heavily occluded, the ground might not be visible and performance degrades. This is a problem for objects in immediate range. A simple yet effective solution is to use grounding only for non-immediate range. Merging random displacement and grounding based on the distance is called *combined* by us, since it combines two ways of augmentation.

The preference of the augmentation method depends on the use case. For best overall performance combined is the method of choice. For short range prediction just random augmentation is preferred and for long range prediction random augmentation with grounding yields best results.

C. Box Prediction

The proposals of Section III.B are used for the final box prediction. They are used in a ROI cropping step on the image and the LiDAR pointcloud. For image ROI cropping, the proposal is projected into the 2D image space. The minimal 2D box containing all corners is then computed. That 2D box is scaled by a factor of 1.5, increasing the context area. Finally, the image is cropped.

The pointcloud gets cropped by using a 3D axis aligned box around the centeroid of the proposal. Similar to RGB, increasing the size of the region crop increases the context available to the neural network. The crop of the pointcloud is then voxelized, using a regular grid of shape 64x64x9. Then, the points in each voxel get encoded as the highest point in a voxel and the density of a voxel.

Both encoded crops are used by two separate 2D convolutional encoders. One operating in BEV and one in image view. As an encoder for images DenseNet [20] works best. Any encoder works, however with different performances, as we show in our ablation study. For the BEV feature representation we use a network architecture inspired by VGG16 [21] explained in our implementation details.

The two encodings are global average pooled separately and then concatenated. The final refinement is then done by a fully connected layer with a ReLU activation and four parallel fully connected layers with different activation functions. The first output regresses a single value with a sigmoid activation function representing the objectness. The second output regresses three values for the position offset $\Delta x, \Delta y, \Delta z$ from the proposal box. The size offsets $\Delta l, \Delta w, \Delta h$ are regressed as another set of three values. The last regression head predicts the orientation as two outputs with tanh activation function. The orientation needs two values for $s_\theta = \sin(\theta)$ and $c_\theta = \cos(\theta)$.

D. Loss

The first component of our loss is for the objectness of an anchor. As the number of positive and negative proposals is unbalanced, focal loss [22] is a good choice. However, we do not know the optimal γ for our use case, therefore, we use the automated focal loss [23] for objectness, defined as

$$L_{objectness} = -(1 - p_t)^{-\log(\hat{p}_t)} \log(p_t), \quad (1)$$

where p_t is the probability of the correct prediction and $\hat{p}_t$ the expected probability of a correct prediction as defined in [23]. Then, for the regression of the centeroid (xyz) and the size (lwh) common smooth l1 loss is used as L_{xyz} and L_{lwh}. The heading θ which is encoded as s_θ, c_θ has a smooth l1 loss on the encoding defined as

$$L_\theta = \texttt{smooth\_l1}\left(\begin{pmatrix} s_\theta \\ c_\theta \end{pmatrix}, \begin{pmatrix} \sin(\theta_{gt}) \\ \cos(\theta_{gt}) \end{pmatrix}\right), \quad (2)$$

where θ is the rotation around the up axis. Finally, all losses are summed, yielding best results without weighing. The total loss for backpropagation is defined as

$$L = L_{xyz} + L_{lwh} + L_\theta + L_{objectness} \quad (3)$$

where L_{xyz}, L_{lwh}, L_θ and $L_{objectness}$ are as defined above.

E. Implementation Details

Our BEV encoder consists of 5 conv-pool-blocks with two 3x3 convolution and batch norm followed by a 2x2 max pooling. The number of filters for the layers are [24, 48], 2x [32, 64], 2x [16, 32]. After that there are two 1x1 convolution layers reducing the feature size to 16 and finally to 10. We chose such low parameter numbers to counter overfitting on the KITTI pedestrian dataset.

LRPD is trained using adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$ and $\epsilon = 10^{-8}$. Our learning rate follows an exponential decay schedule, starting at 10^{-4} and ending at 10^{-6} after 10 epochs of training using a batch size of 32. For regularization we use batch normalization and orthogonal weight initialization inside our network combined with l2 weight decay at a rate of 10^{-3}. The encoders are used in the ImageNet [24] pretrained version provided by `tf.keras`.

For further regularization we use data augmentation. Principal component jitter, left to right flipping, micro translations (0-3px offset) and resolution change (via scaling) are applied on the images. For LiDAR jitter on the xyz points (white noise) is added and a micro translation (0-3cm) is applied to the entire pointcloud. This data augmentation is just distortion of the input and does not replace the proposal augmentation in Section III.B. We apply data augmentation nine times including the augmentation from Section III.B, increasing our training dataset ten fold adding significant variance.

Overall, we have developed LRPD, a specialized approach for long range pedestrian detection, starting with a proposal stage and using RGB instance masks for 3D LiDAR point based proposals. Then, the proposals are refined, predicting the position, size, orientation and objectness in a second stage. To further improve performance proposal augmentation is introduced and automated focal loss is applied.

IV. EVALUATION

We evaluated our LRPD network architecture on the KITTI benchmark focusing on birds-eye-view (BEV), since it is most important for autonomous vehicles with road users moving on the ground plane. The KITTI BEV metric uses top-down IoU to classify correct predictions with a threshold of 0.5 for pedestrian IoU. Furthermore, we inspected performance using precision, recall, f1-score, mean error and mAP with distance from center as a criterion.

A. Results

Table 1 compares LRPD against the top performing open source approaches, reproducing their results from the paper. Our approach is competitive, but KITTI BEV pedestrians benchmark weights samples by occurrence and equally not by distance range. Whereas close objects are over represented in KITTI, the far objects have nearly no importance for the overall score on the KITTI benchmark. However, especially long range detections are important for security and smooth maneuvers. Therefore, we analysed how the performance changes with respect to long range.

.3 Method	Moderate		Easy		Hard	
	all	long range	all	long range	all	long range
AVOD [11]	50.32	12.00	58.49	26.14	46.98	12.09
- VMVS [13]	**50.34**	-	60.34	-	46.45	-
PointPillars [2]	48.64	4.24	57.60	12.14	45.78	4.26
- TANet [16]	51.38	-	60.85	-	47.54	-
F-PointNet [9]	49.57	12.55	57.13	16.30	45.48	10.94
IPOD [12]	49.79	-	**60.88**	-	45.43	-
MMLab-PointRCNN [25]	46.13	2.30	54.77	9.51	42.84	2.62
LRPD w/ Rand. Aug. (ours)	47.04	15.57	51.94	33.92	46.51	15.57
LRPD w/ Grounded (ours)	39.89	**15.63**	45.59	**38.92**	39.88	**15.65**
LRPD w/ Combined (ours)	50.01	15.45	56.41	27.00	**48.46**	13.64

Table 1. The best performing approaches on KITTI pedestrian BEV mAP compared to LRPD (ours). On *all* (ranges) LRPD is comparable to the state-of-the-art, whereas on *long range* (> 30m) LRPD outperforms all reproducible state-of-the-art approaches, not all were reproducible. Since hard examples have slight bias towards distant objects, our LRPD outperforms other state of the art approaches in this category. LRPDs performance can be attributed to it being designed towards long range detection.

We introduce an evaluation metric, that measures performance of object detectors at different ranges, since detection at long range is not separately evaluated in KITTI. To achieve this, we limit the KITTI metric to different intervals. We chose 0-10 meters for the immediate proximity, 10-20 meters for close range, 20-30 meters for medium range and above 30 meters for long range (Figure 4). This means however, that our comparison is limited to publications with source code available. The top performing algorithms on KITTI BEV pedestrian detection are AVOD [11], F-PointNet [9], PointPillars [2] and Point-RCNN [25].

We observed deficits in current state-of-the-art approaches in a good mAP for short range but suffering from worse detection rates at long range. Our LRPD however is able to maintain performance for long range (30+) Figure 4. Even though none of LRPDs configurations can maintain performance perfectly, they outperform the other approaches as can be seen in Figure 2. For an in depth comparison at long range, Table 1 reports our results and the results of AVOD, F-PointNet, PointRCNN and PointPillars for long range. PointPillars which is only LiDAR based with 4.24-12.14 barely detects any object, whereas other approaches leveraging LiDAR and RGB achieve better results (10.94-26.14). Our approach with grounding on the proposal boxes significantly outperforms all others with mAP values of 15.63-38.91. Long range mAP was improved by a factor of 1.5 over the state of the art.

Even though all approaches have similar KITTI BEV mAP values, their failure cases are different. In Figure 5 we show some qualitative results, visualizing the typical failure modes. When two pedestrians are close to each other and overlapping, our approach is able to reliably detect both due to the instance segmentation. LiDAR only approaches have difficulties detecting easily visible objects at long range, leading to PointPillars not predicting long range objects reliably. At short to medium ranges (Figure 5 center column) all approaches perform well with the difference showing on long range detections in the background. However, since our LRPD approach uses person instance masks it confuses cyclists as a pedestrian, like AVOD. F-PointNet and Point-RCNN have very poor precision (0.22 and 0.28) at a maximum recall of 0.62 and 0.55 resulting in the numerous false positives of F-PointNet in Figure 5. In contrast, our approach

Method	F1	Prec.	Rec.	3D Euc. Err.
AVOD [11]	0.48	0.57	0.43	0.14m
PointPillars [2]	0.52	0.66	0.43	**0.11m**
F-PointNet [9]	0.33	0.22	0.62	0.16m
MMLab-PointRCNN [25]	0.37	0.28	0.55	0.12m
LRPD w/ Rand. Aug. (ours)	0.70	0.70	0.69	0.17m
LRPD w/ Grounding (ours)	0.71	0.68	0.74	0.22m
LRPD w/ Combined (ours)	**0.73**	**0.76**	**0.69**	0.18m

Table 2. Precision and recall (at best F1-score) using 1 m euclidean distance as error threshold for assignment

Proposals	Encoder	Refinement Modality	BEV mAP
Baseline	VGG16 [21]	LiDAR + RGB	30.2
Baseline	VGG19 [21]	LiDAR + RGB	39.9
Baseline	ResNet50 [26]	LiDAR + RGB	33.4
Baseline	Xception [27]	LiDAR + RGB	31.3
Baseline	MobileNet [28]	LiDAR + RGB	28.5
Baseline	DenseNet [20]	LiDAR + RGB	43.4
Baseline	DenseNet [20]	LiDAR	30.9
Baseline	DenseNet [20]	RGB	0.3
Rand. Aug.	DenseNet [20]	LiDAR + RGB	51.1
Grounding	DenseNet [20]	LiDAR + RGB	43.5
Combined	*DenseNet [20]*	*LiDAR + RGB*	**54.8**

Table 3. In our ablation studies we show the impact of different design decisions on the performance of our network. The final LRPD configuration is marked in italic.

leads with 0.76 precision at a recall of 0.69 as a maximum. However, regarding 3D euclidean error our network lacks behind other state of the art networks explaining the worse mAP values despite better precision and recall (Table 2). Our detailed quantitative evaluation in Table 1 (long range), Table 2, Figure 2 and Figure 4 is able to capture range based error effects in contrast to the original KITTI metric.

B. Ablation Studies

To verify the effectiveness of all components in our approach, we derived ablation studies. These were performed on the KITTI pedestrian BEV benchmark regarding moderate AP. We swapped out the encoder, changed the input modality and the proposal augmentation separately to observe the impacts. The ablation studies were done on a 80:20 split on the KITTI Benchmark different from the 50:50 split used for our comparison with other approaches to avoid a bias. In our ablation studies (Table 3) we show that both LiDAR and RGB are integral parts to the success of our
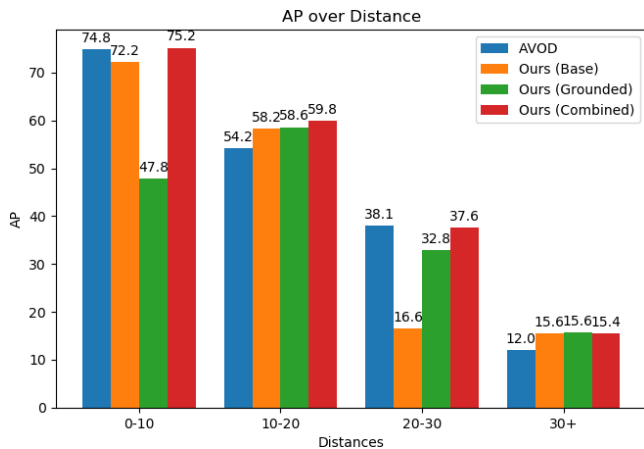
Fig. 4. BEV Detection mAP on KITTI *moderate* for pedestrians plotted for different ranges, using AVOD as the best baseline and our approaches. While AVOD is comparable to the best configuration for immediate (0-10m) to medium range (20-30), LRPD outperforms on long range (30+ m) with all configurations.

approach. By removing RGB or LiDAR from the second stage, described in Section III.C, the performance degrades significantly, highlighting the importance of using both sensor modalities. Furthermore, the encoder choice has as strong impact on performance of the network. Where VGG19 [21] and DenseNet [20] significantly increase the performance, MobileNet [28] slightly decreases the performance.

When varying the proposal method described in Section III.B, we can show that the choice of proposal augmentation has a significant impact on mAP (up to +11.4). Baseline describes creating the proposals without any data augmentation and achieves the worst results of our proposal configurations. Even though random augmentation has a far better overall moderate BEV mAP (Table 3) grounding yields better results for medium and long range detection (Figure 4). Combining the strong performance of random augmentation for short range detection with grounding for medium and long range detection, as described in Section III.B, yields best overall results, with only 0.1 mAP reduction on long range. Despite a low overall BEV mAP of 43.5 grounding achieves the best results on long range (Table 1). This proves that best mAP on KITTI does not correlate to detecting pedestrians better at long range. Therefore, our distance range based analysis of the mAP is mandatory to gain insights into the performance characteristics of approaches, unveiling that on long range detection LRPD significantly outperforms other state-of-the-art approaches.

V. Conclusions

In our analysis, we proved limitations of the state-of-the-art in long range detection performance. Thus, we proposed LRPD leveraging the strengths of LiDAR and RGB for proposal generation and then symmetrically using them in the refinement stage. Our evaluation verifies that our novel approach outperforms the current state-of-the-art significantly for long range 3D detection of pedestrians.

We introduced a new evaluation methodology on KITTI 3D object detection to gain more insights into pedestrian detectors. Using this methodology, we showed that an architecture with a lower overall KITTI BEV score can still be far superior in long range detection than the best performing approaches. Due to the importance of long range detection for comfort and safety of autonomous vehicles, we encourage authors to test their performance not only on the overall metric but to separate the metric into different distance ranges.

REFERENCES

[1] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

[2] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[3] W. Liu, D. Anguelov, D. Erhan, C. Szegedy, S. Reed, C.-Y. Fu, and A. C. Berg, "Ssd: Single shot multibox detector," in *European Conference on Computer Vision (ECCV)*, Springer, 2016.

[4] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[5] T.-Y. Lin, P. Dollár, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[6] R. Girshick, "Fast r-cnn," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.

[7] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014.

[8] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015.

[9] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas, "Frustum pointnets for 3d object detection from rgb-d data," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[10] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[11] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. L. Waslander, "Joint 3d proposal generation and object detection from view aggregation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.

[12] Z. Yang, Y. Sun, S. Liu, X. Shen, and J. Jia, "Ipod: Intensive point-based object detector for point cloud," *Computing Research Repository (CoRR), abs/1812.05276*, 2018.

[13] J. Ku, A. D. Pon, S. Walsh, and S. L. Waslander, "Improving 3d object detection for pedestrians with virtual multi-view synthesis orientation estimation," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[14] Z. Yang, Y. Sun, S. Liu, X. Shen, and J. Jia, "Std: Sparse-to-dense 3d object detector for point cloud," *Computing Research Repository (CoRR), abs/1907.10471*, 2019.

[15] Z. Wang and K. Jia, "Frustum convnet: Sliding frustums to aggregate local point-wise features for amodal 3d object detection," *Computing Research Repository (CoRR), abs/1903.01864*, 2019.

[16] Z. Liu, X. Zhao, T. Huang, R. Hu, Y. Zhou, and X. Bai, "Tanet: Robust 3d object detection from point clouds with triple attention," in *AAAI Conference on Artifical Intelligence*, 2020.

[17] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research (IJRR)*, vol. 32, no. 11, pp. 1231–1237, 2013.
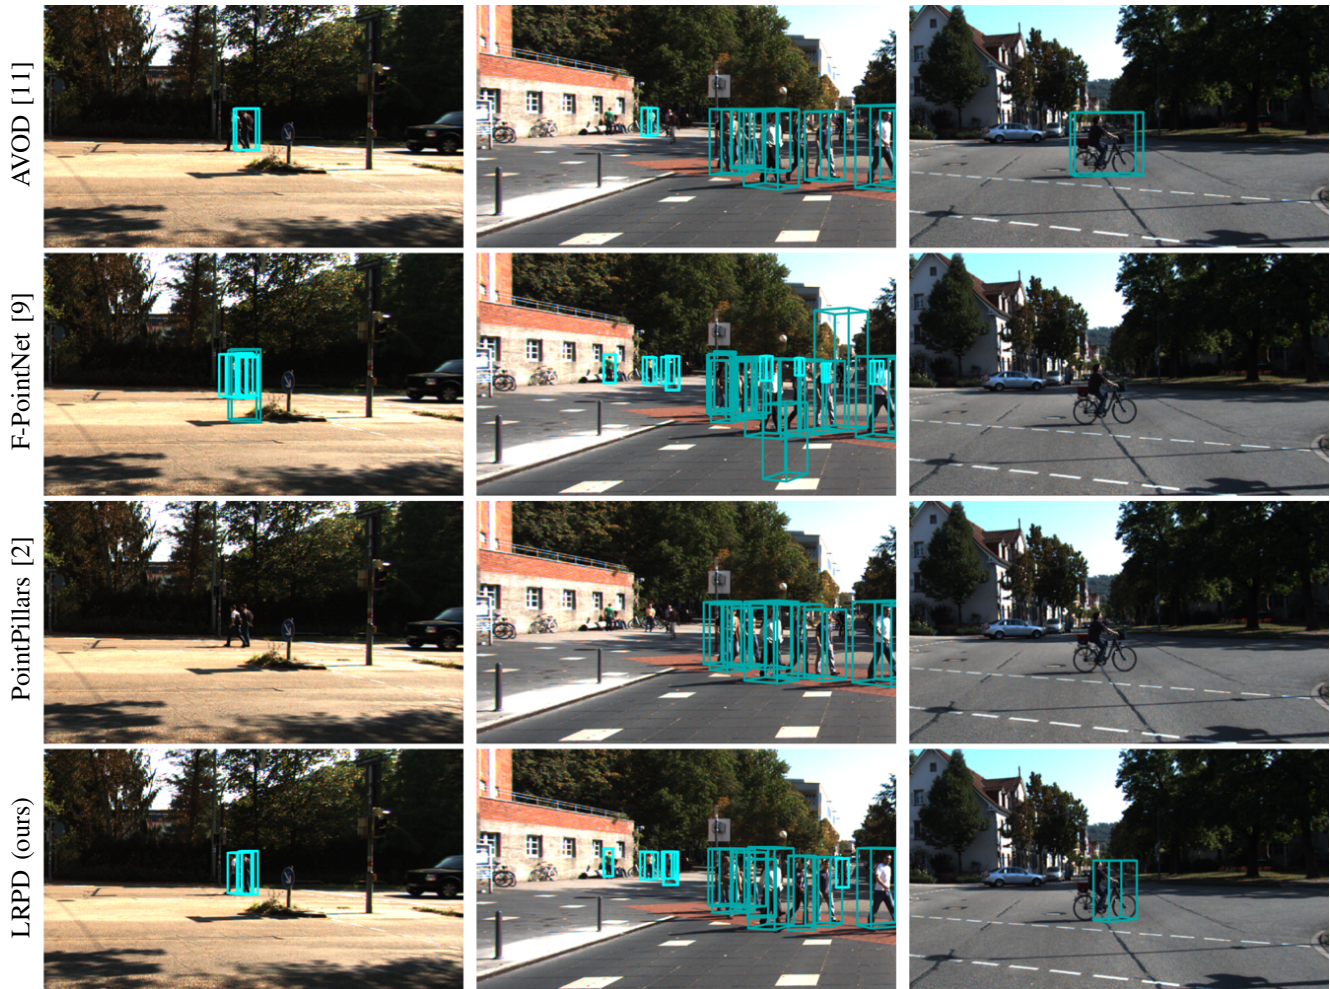
Fig. 5. Qualitative comparison of the state-of-the-art to LRPD (ours). 3D pedestrian detections at less than 20 meters are darker to emphasize on long range detection. The Left column shows two pedestrians. Our Approach and F-PointNet are capable of detecting both, F-PointNet however has false positives (FP). The center column shows a typical prediction, where AVOD and PointPillars are not able to detect far away pedestrians (left) and F-PointNet has a high FP-Rate when still detecting far away objects. Only LRPD is able to detect the far pedestrians while maintaining precision at medium range. Finally, the right column shows a common failure case of pedestrian detectors, AVOD and ours misclassify a person on a bicycle as a pedestrian.

[18] R. Battrawy, R. Schuster, O. Wasenmüller, Q. Rao, and D. Stricker, "Lidar-flow: Dense scene flow estimation from sparse lidar and stereo images," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[19] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[20] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[21] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *Computing Research Repository (CoRR), abs/1409.1556*, 2014.

[22] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal loss for dense object detection," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[23] M. Weber, M. Fürst, and J. M. Zöllner, "Automated focal loss for image based object detection," *Computing Research Repository (CoRR), abs/1904.09048*, 2019.

[24] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.

[25] S. Shi, X. Wang, and H. Li, "Pointrcnn: 3d object proposal generation and detection from point cloud," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.

[26] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[27] F. Chollet, "Xception: Deep learning with depthwise separable convolutions," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

[28] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "Mobilenets: Efficient convolutional neural networks for mobile vision applications," in *Computing Research Repository (CoRR), abs/1704.04861*, 2017.