

Center Feature Fusion: Selective Multi-Sensor Fusion of Center-based Objects

Philip Jacobson¹, Yiyang Zhou¹, Wei Zhan¹, Masayoshi Tomizuka¹ and Ming C. Wu¹

Abstract—Leveraging multi-modal fusion, especially between camera and LiDAR, has become essential for building accurate and robust 3D object detection systems for autonomous vehicles. Until recently, point decorating approaches, in which point clouds are augmented with camera features, have been the dominant approach in the field. However, these approaches fail to utilize the higher resolution images from cameras. Recent works projecting camera features to the bird’s-eye-view (BEV) space for fusion have also been proposed, however they require projecting millions of pixels, most of which only contain background information. In this work, we propose a novel approach Center Feature Fusion (CFF), in which we leverage center-based detection networks in both the camera and LiDAR streams to identify relevant object locations. We then use the center-based detection to identify the locations of pixel features relevant to object locations, a small fraction of the total number in the image. These are then projected and fused in the BEV frame. On the nuScenes dataset, we outperform the LiDAR-only baseline by 4.9% mAP while fusing up to 100x fewer features than other fusion methods.

I. INTRODUCTION

3D object detection is a central pillar of the perception system for modern autonomous vehicles. To tackle this problem, most vehicles employ a full suite of sensors, most prominently cameras and LiDARs. LiDAR, the main workhorse in most detection systems, provides low-resolution but accurate spatial and depth information, whereas camera provides high resolution color and texture information. The process of extracting and combining useful complementary information from each of these sensor modalities, known as sensor fusion, is thus of immense importance to building accurate detection models.

The most recent fusion approaches in the literature can be grouped into two main categories. In the first category, raw LiDAR points are decorated with camera features, such as segmentation class labels or CNN features [1], [2], [3], [4], [5]. In the second, deep CNN features from both modalities are associated and fused together, often using a learnable attention mechanism [6], [7], [8], [9]. Point decoration, while performing well on many detection benchmarks, only supplements already existing LiDAR points, failing to fully leverage the higher camera resolution. Meanwhile, associating deep features from both modalities requires some mechanism to associate features between the two views, which is a challenging problem in and of itself.

Recently, fusion in the bird’s-eye-view representation has emerged as a compelling approach. BEVFusion [9] has

recently achieved state-of-the-art performance on both the nuScenes [10] and Waymo Open [11] 3D object detection challenges. BEV, unlike the camera’s perspective view, is an attractive representation for fusion since it preserves the geometry of the scene. However, projecting camera pixels to BEV is much less efficient than projecting 3D points to perspective view because a) we require per-pixel depth estimates to perform the projection and b) the number of camera pixels far outnumbers the number of points in a LiDAR pointcloud.

In this paper, we propose a novel fusion approach which also leverages the power of the BEV representation. However, we argue that projecting all of the camera features into BEV is wasteful, since the vast majority of pixels simply represent background. Instead, we leverage center-based detector CenterNet [12] to identify objects in perspective view. Since CenterNet infers bounding box characteristics from only a single point, all of the relevant features are agglomerated in a single location per object, making projection to BEV natural and efficient. We therefore set a confidence threshold in the detection heatmap which is used to decide which pixels to project to BEV. After these features are projected, the BEV features pass through a fully-convolutional encoder to help with alignment. Lastly, the aligned features are concatenated with BEV LiDAR features before a CenterPoint regression head calculates the final bounding box locations.

Additionally, the issue of preserving alignment when applying multi-modal data augmentation is essential in fusion work. To enable the training of our model, we apply data augmentation in a novel fashion: first, we apply standard LiDAR augmentations (random flip, scaling, rotation, and translation) [13]. Next, we save the random parameters of the augmentation. Once the camera features are projected into global 3D coordinates, we then apply the augmentation on the camera features. This allows us to apply the exact same augmentations to both set of features in an efficient manner.

In summary, our contributions are the following:

- 1) We explore a selective deep feature fusion strategy in the bird’s-eye-view domain through projecting camera features based on their confidence score in a detection heatmap.
- 2) We propose a novel mechanism for data augmentation that aligns features projected to bird’s-eye-view.
- 3) We validate our method on the nuScenes dataset, achieving a significant +4.9% improvement in Mean Average Precision (mAP) and +2.4% nuScenes Detec-

¹University of California, Berkeley {philip_jacobson, yiyang.zhou, wzhan, tomizuka}@berkeley.edu
wu@eecs.berkeley.edu

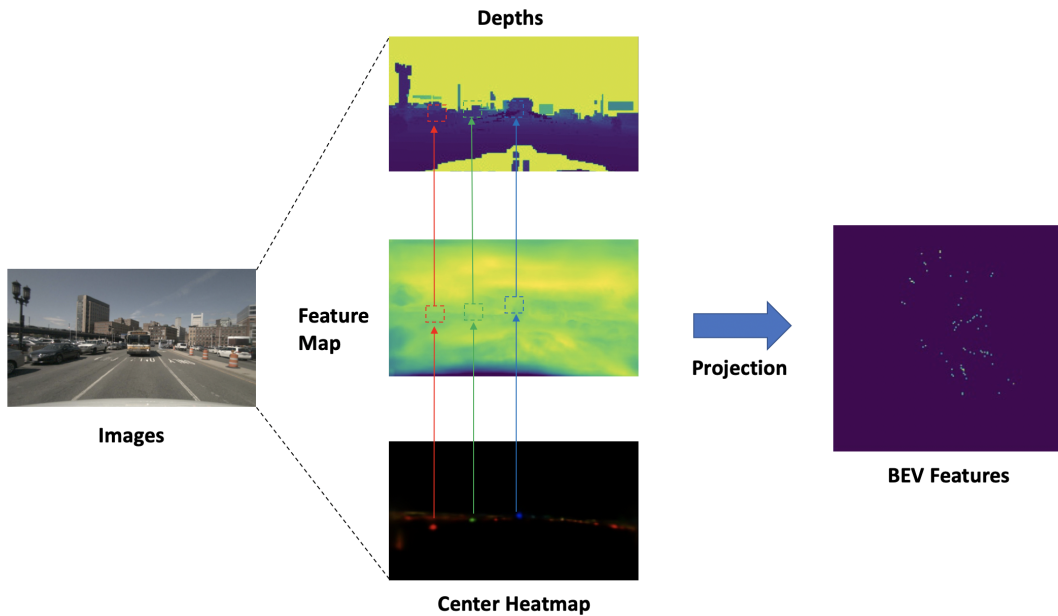


Fig. 1: Our proposed selective feature projection. The detection heatmap generated by CenterNet’s detection head is used to select interesting regions from the corresponding deep feature map. Dense depth predictions are then used to project the corresponding features to BEV, where they can be fused with LiDAR BEV features.

tion Score (NDS) over a LiDAR-only baseline. Our performance is in-line with other fusion methods, even as we fuse a far smaller number of features.

II. RELATED WORK

A. LiDAR-based 3D Object Detection

LiDAR-based 3D object detectors aim to detect objects with a given point cloud, represented as an unordered set of 3D points in space. PointNet and its derivatives [14], [15] operate directly on said point set. Voxel-based methods first quantize point clouds into regular 3D voxels before the voxelized representation is processed through the detector stage. VoxelNet [16] discretizes the 3D space before processing points inside each voxel with a PointNet-like network; voxels are then processed using a dense 3D convolutional neural network. SECOND [17] improves on VoxelNet by introducing efficient sparse convolutions. Pillar-based methods discretize the point cloud into vertical columns, flattening the representation along the height dimension. PIXOR [18] flattens the point cloud into a 2D pseudo-image which is encoded with a 2D CNN. PointPillars [19], similar to VoxelNet, processes points inside pillars through a PointNet before the encoded pillars are passed through a 2D convolutional network. CenterPoint replaced the standard two-stage detection architecture with an anchor-free approach by utilizing a center-based representation of detected objects, achieving the new state-of-the-art on 3D detection tasks [20].

B. Camera-based 3D Object Detection

3D object detection using only camera images, although attractive as an alternative to expensive 3D sensors, still lags

significantly behind LiDAR-based detection due to the challenge of monocular depth estimation. Pseudo-lidar methods use a pre-trained depth estimation network to transform RGB images into point clouds that are then compatible with standard LiDAR-based detectors [21], [22]. Another class of detectors transforms camera features into BEV representation. BEVFormer [23] achieves state-of-the-art results by utilizing a temporal cross-attention to calculate interactions between camera features in both the current and previous frame. CenterNet, the 2D counterpart to CenterPoint, similarly uses center-based object detection in perspective view combined with a depth-estimation head to generate 3D bounding box characteristics [12].

C. Multi-modal Fusion

Because of the unique advantages and disadvantages of different sensor modalities, leveraging data from each of these modalities is the most promising approach for 3D object detection. Early sensor fusion work focused primarily on proposal-level fusion, i.e. fusing information at the region proposal level. MV3D [24] fuses object proposals in three views (BEV, front view, and perspective view), whereas AVOD fuses them in two views (BEV and perspective view) [25]. RoarNet determines geometrically feasible 3D poses from images to generate a set of 3D region proposals [26]. Frustum-PointNet [27] and Frustum-ConvNet [28] lift image proposals into 3D frustums to identify relevant regions of the point cloud. FUTR3D [7] and TransFusion [8] generate object proposals in 3D and refine them using a transformer decoder which attends to 2D features.

Point decoration has been a popular paradigm in sensor fusion for the past few years. PointPainting [1] proposed

projecting the point cloud onto the outputs of 2D semantic segmentation networks to augment each point with the class label of the corresponding 2D pixel. FusionPainting [3] fuses segmentation labels from both 2D and 3D segmentations to generate more accurate labels for the point cloud. PointAugmenting [2] instead annotates points with deep CNN features while processing the LiDAR and camera features through separate streams in the CNN backbone. AutoAlign [4] and AutoAlignV2 [5] use a learnable cross-attention module to associate pixel-level features with voxels, instead of the hard association of projection using a camera matrix. MVP [29] takes a different input decorating approach, projecting the point cloud onto a segmentation map to generate virtual LiDAR points inside each object to densify the point cloud.

Most recently, mid-level fusion in which fusion is done at the deep-feature level has become popular. Continuous Fusion [30] shares information between the 2D and 3D backbones at all layers of the neural network. DeepFusion [6] uses a learnable attention layer to associate CNN features between different sensor modalities. BEVFusion [9] uses the approach from LSS [31] to generate per-pixel depth predictions to project the full set of camera features to BEV, where after undergoing a pooling operation, they are concatenated with LiDAR BEV features.

III. CENTER FEATURE FUSION

This section introduces our proposed method, Center Feature Fusion. Our model uses CenterPoint [20] and CenterNet [12] as 3D and 2D feature encoders, respectively, to generate features for fusion. The CenterNet encoder additionally generates detection heatmaps in the perspective view, which we use to select the relevant features for projection. We generate per-pixel depth estimates and then use these to project the selected camera features to BEV. The camera BEV features are encoded using a further 2D convolutional network before concatenation with LiDAR features. The concatenated features are then passed to CenterPoint’s various bounding box heads. A detailed diagram illustrating this method is shown in Fig. 2. Additionally, we introduce a novel data augmentation scheme, ProjectionAlign, to facilitate the training of the fusion model. To ensure alignment between 2D and 3D features, we apply data augmentation to the projected 2D camera features, and not the images at the input, so that the same exact augmentations can be applied to both set of features.

A. Camera Feature Representation

Given the high resolution of cameras relative to LiDARs, the 2D feature extractors of most fusion models generate significantly more features than those of the LiDAR backbone. However, utilizing all of these features can represent a significant bottleneck; projecting potentially millions of pixels to BEV is a computationally intensive task, while allowing each LiDAR feature to attend to all camera features greatly increases model complexity. We observe that the vast majority of camera features are irrelevant for object detection, only representing the background. CenterNet takes

this to the extreme by relying on only the features located at a single center pixel per object to characterize it. Thus, we argue for an approach which selectively chooses relevant features to fuse, utilizing CenterNet’s 2D detection heatmap.

B. Camera Detection Pipeline

Before fusion, camera images are fed into a pre-trained CenterNet object detector. CenterNet’s heatmap head generates a keypoint heatmap $\hat{Y} \in [0, 1]^{\frac{w}{s} \times \frac{h}{s} \times K}$, where w and h are the input height and width, respectively, s is the downsampling factor of the output, and K is the number of classes. Each local peak in the heatmap (i.e. a point with value larger than its eight neighbors) corresponds to a detected object, and the bounding box characteristics are then inferred from the features located at that point. Because of this, we argue that all of the information relevant for detection and bounding box regression are located at these central pixels. We selectively fetch these points by choosing a heatmap threshold $\in [0, 1]$, pixels below the threshold are disregarded. A lower threshold, while fetching more features per object as well as features for less confident detections, increases latency. For each pixel location above this threshold, the corresponding location in the output feature map of the DLA34 [32] CenterNet backbone is chosen as the camera feature. Fig. 1 illustrates our selective projection method.

In order to project the selected features from perspective view to BEV, we require per-pixel depth estimates. While BEVFusion utilizes the approach in LSS of predicting a per-pixel depth distribution [9], this requires training a separate view transformer. Instead, we project LiDAR points to the camera frame and then perform depth completion. MVP utilizes a nearest-neighbor based classifier for depth completion to generate the virtual points, the simplest approach to depth completion [29]. Many deep-learning approaches exist for depth-completion [33], [34], [35], [36], however these require separate training and are computationally expensive. Instead, we use IP-Basic [37], a fast, lightweight and efficient classical depth completion algorithm which is significantly more accurate than a nearest-neighbor classifier. To prevent objects in the background outside of the LiDAR range from being projected into the foreground, we only perform depth completion on points within interpolation range, and set all other depths outside of the detection range. Using the per-pixel depths estimates, we project the previously selected perspective view features to the 3D world coordinate system. Features are then aggregated inside the cells of the BEV grid, and a max pooling operation is performed on all of the features within a cell to produce the final BEV features.

C. Fusion in BEV Frame

Once the camera features are organized into the 2D BEV grid, we further encode them using a small 2D convolutional neural network, as in PointAugmenting and BEVFusion. This extra convolutional stream allows the model to compensate for small misalignments from the inherent inaccuracies in the depth estimation process. After this additional refinement, the camera BEV features are concatenated with the LiDAR

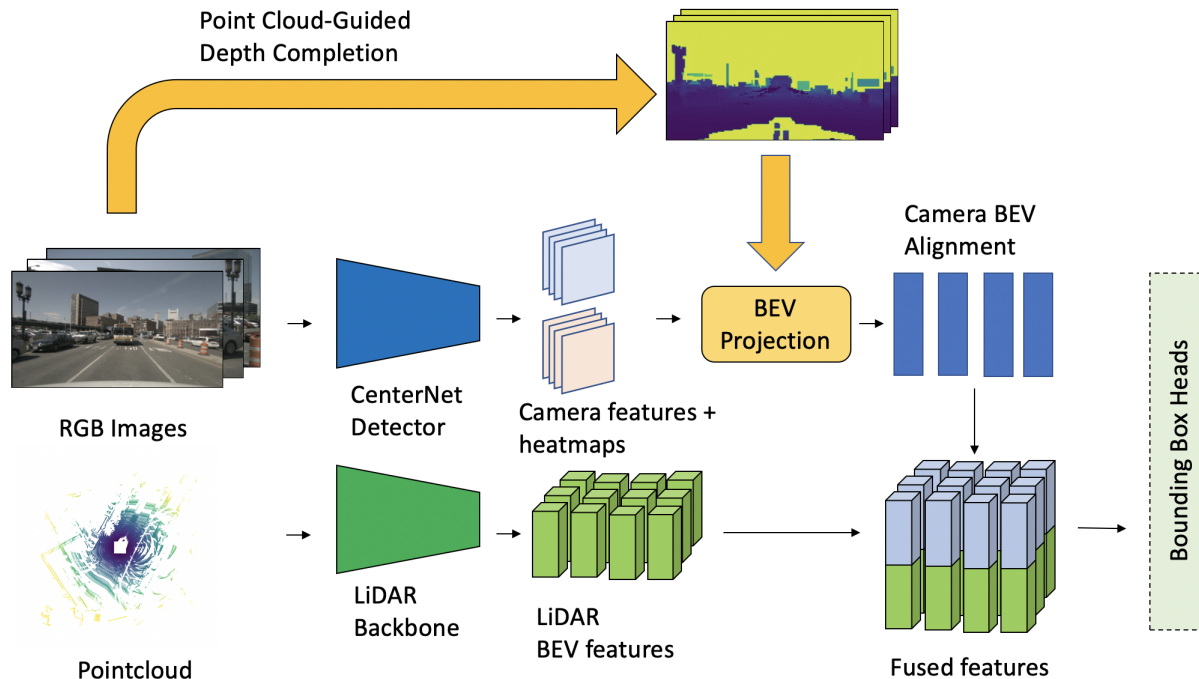


Fig. 2: Summary of our proposed Center Feature Fusion model. The pointcloud and camera images are first encoded using a VoxelNet and DLA34-CenterNet, respectively. Our selective projection is performed on the camera features to project them to BEV where they are then fed into a small 2D CNN to help with feature alignment. The multi-modal BEV features are then concatenated and passed into CenterPoint’s bounding box heads.

features, where they then undergo a shared convolutional block before being broadcast to CenterPoint’s regression heads.

D. ProjectionAlign Augmentation

Strong data augmentations are essential during training for most current 3D detection benchmarks to achieve state-of-the-art results. In sensor fusion applications, data augmentation poses a challenge due to the fundamentally different approaches applied in 2D (e.g. random cropping, random flipping) and 3D (random rotation, scaling). Applying these augmentations naively destroys the consistency between fused points and degrades performance. DeepFusion addresses this problem with InverseAug, in which they save the 3D augmentation parameters, reverse them at the fusion stage, and then calculate the feature’s corresponding coordinates in the 2D frame. By using BEV as our fusion representation, we are able to propose an even simpler algorithm we call ProjectionAlign. Because we project camera features to BEV, we generate a pseudo-3D point cloud, which can be augmented in the same manner as the LiDAR point cloud. Thus, we simply apply augmentations to the point cloud, save the exact parameters of the augmentation, and apply them to the camera features once they are converted to the same 3D world coordinate system. We don’t apply any geometric augmentations to the camera images directly to preserve the exact alignment between the two modalities. A visualization of our alignment scheme is shown in Fig. 3.

IV. EXPERIMENTS

We evaluate CFF on the nuScenes dataset [10], a large-scale autonomous driving benchmark. nuScenes consists of 1000 scenes, split into 700 for training, 150 for validation, and 150 for testing. Object annotations are included for every tenth frame in the dataset, resulting in 28130 frames for training, 6019 frames for validation, and 6008 frames for testing. Each frame consists of a 32-beam LiDAR scan and six monocular camera images of resolution 1600 x 900 covering the full 360-degree field-of-view. Annotations come from a set of 10 classes. Detection performance is measured using two main metrics: Mean Average Precision (mAP), and the nuScenes Detection Score (NDS).

A. Implementation Details

We implement CFF with the open-source code of CenterPoint and CenterNet. The CenterNet model uses a DLA34 backbone which is pre-trained separately on the nuScenes dataset. Our CenterPoint model is built with a VoxelNet backbone. Following CenterPoint, we use a detection range of $[-54m, 54m]$ for the X and Y axis, and a range of $[-5m, 3m]$ for the Z axis. We use a voxel size of $(0.075m, 0.075m, 0.2m)$. We perform the standard LiDAR data augmentations of random flipping along the X and Y axis, global random scaling, global random rotation, and global random translation. We use the adamW optimizer [38] with a one-cycle learning rate policy, a max learning rate of

	mAP	NDS	Car	Truck	C.V.	Bus	Trailer	Barrier	Motor.	Bike	Ped.	T.C.
CenterPoint [20]	59.6	66.8	85.5	58.6	17.1	71.5	37.2	68.5	58.9	43.3	85.1	69.7
Ours	64.5	69.2	85.5	59.1	22.5	73.4	43.2	68.8	72.8	62.1	86.4	70.6
Improvement	+4.9	+2.4	+0.0	+0.5	+5.4	+1.9	+6.0	+0.3	+13.9	+18.8	+1.3	+0.9

TABLE I: Performance comparison of CenterPoint and our method on nuScenes validation set. We report the total Mean Average Precision (mAP) and nuScenes Detection Score (NDS), and per-class mAP. C.V., Motor., Ped., and T.C. are short for Construction Vehicle, Motorcycle, Pedestrian, and Traffic Cone, respectively.

	# of fused features	mAP	NDS	Car	Truck	C.V.	Bus	Trailer	Barrier	Motor.	Bike	Ped.	T.C.
CenterPoint [20]	-	60.3	67.3	85.2	53.5	20.0	63.6	56.0	71.1	59.5	30.7	84.6	78.4
PointAugmenting [2]	~200000	66.8	71.0	87.5	57.3	28.0	65.2	60.7	72.6	74.3	50.9	87.9	83.6
TransFusion [8]	~100000	68.9	71.7	87.1	60.0	33.1	68.3	60.8	78.1	73.6	52.9	88.4	86.7
BEVFusion [9]	~100000	70.2	72.9	88.6	60.1	39.3	69.8	63.8	80.0	74.1	51.0	89.2	86.5
Ours	~ 1000	65.2	69.9	84.7	55.4	26.0	66.2	59.2	74.3	72.6	50.7	85.7	77.5

TABLE II: Comparison with state-of-the-art fusion methods and CenterPoint on nuScenes test set. Our method achieves a similar increase in performance to other fusion approaches, even as we use only approximately 1/100 the number of features for fusion.

$1e-3$, weight decay of 0.1 and momentum between 0.85 and 0.95. We adopt the class-balanced sampling and class-grouped training proposed in CBGS [13]. We train our model using a batch size of 12 on 3 Nvidia A6000 GPUs.

Augmentation Strategy	mAP	NDS
No Augmentation	60.4	66.6
+ LiDAR Only	63.2	68.2
+ ProjectionAlign	64.5	69.2

TABLE III: Ablation study of ProjectionAlign. Results shown for nuScenes validation set.

B. Main Results

First we compare the improvement from our proposed fusion approach against the CenterPoint LiDAR-only baseline. We display this comparison on the nuScenes validation set in Table I. Our method achieves a +4.9% improvement in mAP and a +2.4% improvement in NDS overall against CenterPoint, with improvements in mAP seen across all of the classes. Fig 4 visualizes our results qualitatively for an example frame in the nuScenes validation set.

Table II compares the performance of our method with other state-of-the-art fusion approaches on the nuScenes test set. In addition to mAP and NDS, we also compare an approximate number of fused features per-frame. PointAugmenting [2] (and other related point-decoration methods) augment each point in the point cloud, which is, on average, around 200000 points in nuScenes after concatenation with non-key frames. TransFusion [8] calculates attention between LiDAR BEV features and CenterNet [12] camera features in perspective view, and so the fused features are aggregated across all more than 100000 pixels in the CenterNet output feature maps. BEVFusion [9] projects and fuses all camera features encoded by a Swin-T backbone, amounting to almost 150000 total pixels. In comparison, through our




Fig. 3: Demonstration of ProjectionAlign. a) GT detections in BEV. b) Unaligned projected camera features. c) Aligned camera features. Much stronger location correspondence exists between aligned features and ground truth object locations vs. unaligned features.

selective fusion approach, we toss out the majority of camera features and fetch only around 1000 features to project and fuse, more than $100\times$ fewer than the aforementioned fusion methods. Despite this, we achieve comparable improvements in performance over the CenterPoint baseline. We outperform PointAugmenting on a handful of classes (e.g. bus, barrier) even without the test-time-augmentation PointAugmenting employs. On the more fine-grained classes of motorcyclists and bicyclist, which rely heavily on semantic information




Fig. 4: Qualitative Comparison of detection results from CenterPoint (left) and Center Feature Fusion (right). Red boxes indicate ground-truth objects, while blue indicate model predictions. Our approach detects several cars and trucks that CenterPoint misses, as well as generating more accurate boxes for a number of objects.

for detection, we achieve performance in-line with all of the other methods.

Threshold	# of projected pixels	Latency	mAP
0.5	110	57 ms	60.8
0.1	1442	74 ms	64.5
0.05	2855	98 ms	64.8
0.01	20077	558 ms	63.5
0.0	134400	3690 ms	-

TABLE IV: Comparison of average projection latency, number of pixels projected from perspective view to BEV and resulting mAP for various detection heatmap thresholds on nuScenes dataset. Projecting all camera pixels presents a significant processing bottleneck. We omit the mAP for a threshold of 0.0 as the latency makes training this model infeasible.

C. Ablation Studies

Augmentation Strategy: We validate the effectiveness of our ProjectionAlign data augmentation by comparing its performance against no augmentation, augmentation only the standard LiDAR augmentations are applied (i.e. camera and LiDAR features are misaligned), and our augmentation strategy. The results are summarized in Table III. Without any sort of data augmentation, our fusion model quickly runs into overfitting and improvement over the LiDAR-only baseline is minimal at +0.8 mAP. Applying the naive augmentation strategy improves our model by significant amount of +2.8 mAP. By introducing our aligned augmentation strategy, the model performance is further improved by +1.2 mAP. This confirms the effectiveness of our augmentation scheme and the general importance of aligning fused features, as was shown in [6].

Heatmap Thresholding: To illustrate the computational efficiency gained from our selective projection strategy, we perform experiments varying the heatmap threshold used to determine which camera features should be projected to BEV. These are summarized in Table IV. A threshold of 0, meaning every pixel across all 6 cameras is projected, requires on average more than 3.5 seconds on a single NVIDIA A6000 GPU. Given that most of the pixels are identified by the object detector as background, even a small threshold of 0.01 reduces the number of projected pixels and the required latency by more than $6\times$. We find that improvement in mAP levels off at a threshold of 0.1, which results in fusing almost 1/100 of the total camera features available, illustrating the redundancy of the vast majority of available features. Thus, we use the threshold of 0.1 in all of the experiments presented in this paper.

V. CONCLUSION

In this paper, we present a novel fusion approach, Center Feature Fusion, and an associated novel data augmentation scheme, ProjectionAlign. We demonstrate that working in the BEV feature space, fusing only a small number of camera features significantly improves performance of the LiDAR-only baseline, and CFF is comparable in performance to other fusion methods which requires fusing and correlating many times more features. These results show that a more selective approach in choosing which deep features to fuse may allow for improvements in processing speed and computational efficiency without significantly degrading detection accuracy.

ACKNOWLEDGMENT

We would like to thank Aaron Drumm for his helpful discussions during the writing of the paper. Philip Jacobson is supported by a National Defense Science and Engineering Graduate Fellowship.

REFERENCES

- [1] S. Vora, A. H. Lang, B. Helou, and O. Beijbom, "Pointpainting: Sequential fusion for 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.
- [2] C. Wang, C. Ma, M. Zhu, and X. Yang, "Pointaugmenting: Cross-modal augmentation for 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 11 794–11 803.
- [3] S. Xu, D. Zhou, J. Fang, J. Yin, Z. Bin, and L. Zhang, "Fusionpainting: Multimodal fusion with adaptive attention for 3d object detection," in *ITSC*, 2021.
- [4] Z. Chen, Z. Li, S. Zhang, L. Fang, Q. Jiang, F. Zhao, B. Zhou, and H. Zhao, "Autoalign: Pixel-instance feature aggregation for multi-modal 3d object detection," *arXiv preprint arXiv:2201.06493*, 2022.
- [5] Z. Chen, Z. Li, S. Zhang, L. Fang, Q. Jiang, and F. Zhao, "Autoalignv2: Deformable feature aggregation for dynamic multi-modal 3d object detection," *arXiv preprint arXiv:2207.10316*, 2022.
- [6] Y. Li, A. W. Yu, T. Meng, B. Caine, J. Ngiam, D. Peng, J. Shen, B. Wu, Y. Lu, D. Zhou, Q. V. Le, A. Yuille, and M. Tan, "Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection," in *CVPR*, 2022.
- [7] X. Chen, T. Zhang, Y. Wang, Y. Wang, and H. Zhao, "Futr3d: A unified sensor fusion framework for 3d detection," *arXiv preprint arXiv:2203.10642*, 2022.
- [8] X. Bai, Z. Hu, X. Zhu, Q. Huang, Y. Chen, H. Fu, and C.-L. Tai, "TransFusion: Robust Lidar-Camera Fusion for 3d Object Detection with Transformers," *CVPR*, 2022.
- [9] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "Befusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," *arXiv*, 2022.
- [10] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," *arXiv preprint arXiv:1903.11027*, 2019.
- [11] P. Sun, H. Kretschmar, X. Dotiwalla, A. Chouard, V. Patnaik, P. Tsui, J. Guo, Y. Zhou, Y. Chai, B. Caine, V. Vasudevan, W. Han, J. Ngiam, H. Zhao, A. Timofeev, S. Ettinger, M. Krivokon, A. Gao, A. Joshi, S. Zhao, S. Cheng, Y. Zhang, J. Shlens, Z. Chen, and D. Anguelov, "Scalability in perception for autonomous driving: Waymo open dataset," in *CVPR*, 2020.
- [12] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," in *arXiv preprint arXiv:1904.07850*, 2019.
- [13] B. Zhu, Z. Jiang, X. Zhou, Z. Li, and G. Yu, "Class-balanced grouping and sampling for point cloud 3d object detection," *arXiv preprint arXiv:1908.09492*, 2019.
- [14] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "Pointnet: Deep learning on point sets for 3d classification and segmentation," *arXiv preprint arXiv:1612.00593*, 2016.
- [15] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "Pointnet++: Deep hierarchical feature learning on point sets in a metric space," *arXiv preprint arXiv:1706.02413*, 2017.
- [16] Y. Zhou and O. Tuzel, "Voxelnet: End-to-end learning for point cloud based 3d object detection," in *CVPR*, 2018.
- [17] Y. Yan, Y. Mao, and B. Li, "Second: Sparsely embedded convolutional detection," *Sensors*, vol. 18, no. 10, 2018. [Online]. Available: <https://www.mdpi.com/1424-8220/18/10/3337>
- [18] B. Yang, W. Luo, and R. Urtasun, "Pixor: Real-time 3d object detection from point clouds."
- [19] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, "Pointpillars: Fast encoders for object detection from point clouds," *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12 689–12 697, 2019.
- [20] T. Yin, X. Zhou, and P. Krähenbühl, "Center-based 3d object detection and tracking," *CVPR*, 2021.
- [21] Y. Wang, W.-L. Chao, D. Garg, B. Hariharan, M. Campbell, and K. Weinberger, "Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving," in *CVPR*, 2019.
- [22] R. Qian, D. Garg, Y. Wang, Y. You, S. Belongie, B. Hariharan, M. Campbell, K. Q. Weinberger, and W.-L. Chao, "End-to-end pseudo-lidar for image-based 3d object detection," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5881–5890.
- [23] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, "Bevformer: Learning bird's-eye-view representation from multi-camera images via spatiotemporal transformers," *arXiv preprint arXiv:2203.17270*, 2022.
- [24] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3d object detection network for autonomous driving," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, jul 2017, pp. 6526–6534. [Online]. Available: <https://doi.ieeecomputersociety.org/10.1109/CVPR.2017.691>
- [25] J. Ku, M. Mozifian, J. Lee, A. Harakeh, and S. Waslander, "Joint 3D proposal generation and object detection from view aggregation," *IROS*, 2018.
- [26] K. Shin, Y. P. Kwon, and M. Tomizuka, "Roarnet: A robust 3d object detection based on region approximation refinement," in *2019 IEEE Intelligent Vehicles Symposium (IV)*, 2019, pp. 2510–2515.
- [27] C. R. Qi, W. Liu, C. Wu, H. Su, and L. J. Guibas, "Frustrum pointnets for 3d object detection from rgb-d data," *arXiv preprint arXiv:1711.08488*, 2017.
- [28] Z. Wang and K. Jia, "Frustrum convnet: Sliding frustrums to aggregate local point-wise features for amodal 3d object detection," in *2019 IEEE/RSSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2019, pp. 1742–1749.
- [29] T. Yin, X. Zhou, and P. Krähenbühl, "Multimodal virtual point 3d detection," *NeurIPS*, 2021.
- [30] M. Liang, B. Yang, S. Wang, and R. Urtasun, "Deep continuous fusion for multi-sensor 3d object detection," in *ECCV*, 2018.
- [31] J. Philion and S. Fidler, "Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d," in *Proceedings of the European Conference on Computer Vision*, 2020.
- [32] F. Yu, D. Wang, E. Shelhamer, and T. Darrell, "Deep layer aggregation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [33] M. Jaritz, R. D. Charette, E. Wirbel, X. Perrotton, and F. Nashashibi, "Sparse and dense data with cnns: Depth completion and semantic segmentation," in *2018 International Conference on 3D Vision (3DV)*, 2018, pp. 52–60.
- [34] W. Van Gansbeke, D. Neven, B. De Brabandere, and L. Van Gool, "Sparse and noisy lidar completion with rgb guidance and uncertainty," in *2019 16th International Conference on Machine Vision Applications (MVA)*, 2019, pp. 1–6.
- [35] A. Eldesokey, M. Felsberg, and F. S. Khan, "Confidence propagation through cnns for guided sparse depth regression," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 42, no. 10, pp. 2423–2436, 2020.
- [36] Y. Xu, X. Zhu, J. Shi, G. Zhang, H. Bao, and H. Li, "Depth completion from sparse lidar data with depth-normal constraints," in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [37] J. Ku, A. Harakeh, and S. L. Waslander, "In defense of classical image processing: Fast depth completion on the cpu," in *2018 15th Conference on Computer and Robot Vision (CRV)*. IEEE, 2018, pp. 16–22.
- [38] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.