

AutoAlign: Pixel-Instance Feature Aggregation for Multi-Modal 3D Object Detection

Zehui Chen¹, Zhenyu Li², Shiquan Zhang³, Liangji Fang³, Qinghong Jiang³,
Feng Zhao^{*1}, Bolei Zhou⁴, Hang Zhao⁵

¹ University of Science and Technology, ² Harbin Institute of Technology
³ SenseTime Research, ⁴ The Chinese University of Hong Kong, ⁵ IIS, Tsinghua University

Abstract

Object detection through either RGB images or the LiDAR point clouds has been extensively explored in autonomous driving. However, it remains challenging to make these two data sources complementary and beneficial to each other. In this paper, we propose *AutoAlign*, an automatic feature fusion strategy for 3D object detection. Instead of establishing deterministic correspondence with camera projection matrix, we model the mapping relationship between the image and point clouds with a learnable alignment map. This map enables our model to automate the alignment of non-homogenous features in a dynamic and data-driven manner. Specifically, a cross-attention feature alignment module is devised to adaptively aggregate *pixel-level* image features for each voxel. To enhance the semantic consistency during feature alignment, we also design a self-supervised cross-modal feature interaction module, through which the model can learn feature aggregation with *instance-level* feature guidance. Extensive experimental results show that our approach can lead to 2.3 mAP and 7.0 mAP improvements on the KITTI and nuScenes datasets, respectively. Notably, our best model reaches 70.9 NDS on the nuScenes testing leaderboard, achieving competitive performance among various state-of-the-arts.

1 Introduction

Recent advances in deep learning bring rapid progress in autonomous driving. 3D object detection through LiDAR points plays an essential role in understanding the surroundings of the vehicle. LiDAR points can capture precise 3D spatial information for object detection, however, they often suffer from both the lack of semantic information and sparsity of reflected points, leading to failures under foggy or crowded circumstances. Compared to point clouds, RGB images have better strength in providing semantic and long-distance information. Hence, many approaches explore the data fusion of RGB camera and LiDAR sensors to improve the performance of 3D object detection.

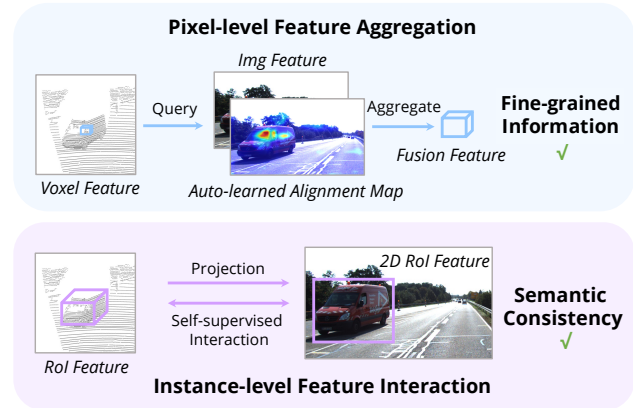


Figure 1: The illustration of interactions between image and point cloud in AutoAlign. The interactions act on two levels: (i) the pixel-level feature aggregation preserves the fine-grained RGB features from images and (ii) the instance-level feature interaction enhances the semantic consistency between non-homogenous representations.

Multi-modal 3D object detectors can be roughly grouped into two categories: decision-level fusion and feature-level fusion. The former detects objects in respective modalities and then ensembles the boxes together in the 3D space [Pang *et al.*, 2020]. Different from the decision-level fusion, the feature-level fusion combines the multi-modal features into a single representation from which the objects are detected. Therefore, the detector can fully utilize features from different modalities at the inference stage. In light of this, more approaches for feature-level fusion have been recently developed. One line of work [Vora *et al.*, 2020; Xie *et al.*, 2020] projects each point to the image plane and gets the corresponding image features through bilinear interpolation. Although the feature aggregation is performed delicately at the *pixel level*, in doing so we will lose the dense patterns in the image domain due to the sparsity of the fusion points, *i.e.*, breaking the *semantic consistency* in the image features. Another line of work [Chen *et al.*, 2017] uses initial proposals provided by 3D detectors to obtain respective RoI features at different modalities and concatenates them together for feature fusion. It maintains the semantic consistency by conducting *instance-level* fusion, however, it suffers from the coarse feature aggregation and the absence of 2D information at the initial proposal generation phase.

To take the best of these two types of approaches, we propose an integrated multi-modal feature fusion framework for 3D object detection, named **AutoAlign**. It enables the detector to aggregate cross-modal features in an adaptive way, which proves to be effective in modeling relationships between non-homogenous representations. Meanwhile, it takes advantage of the fine-grained feature aggregation at the pixel level, but at the same time preserves the semantic consistency through instance-wise feature interaction (see Figure 1).

Specifically, to preserve concrete details in the RGB data, we design a Cross-Attention Feature Alignment (CAFA) module, which dynamically attends the pixel-level features from the image and maintains the efficiency by fusing features at a higher 3D level (pillar or voxel). Each voxel feature will query the whole image plane to get a pixel-wise semantic alignment map. Then, CAFA aggregates the image features based on the alignment map and concatenates them together with the original 3D features. In order to facilitate the learning of semantic consistency between point clouds and images, we propose a novel Self-supervised Cross-modal Feature Interaction (SCFI) module. In detail, we first use paired 2D-3D proposals, predicted by the detector, to extract regional features in their respective domain. After that, a similarity loss will be exerted between paired regional features in 2D and 3D space. By interacting with cross-modal features at the instance level, SCFI strengthens the ability to perceive semantic-related information in CAFA.

Additionally, inspired by multi-task learning, we devise a 2D-3D detection joint training paradigm to regularize the optimization of image branch. Such a training scheme prevents the overfitting issue of the image backbone and further enhances the performance of 3D detectors.

The main contributions of this work are three-fold:

- We propose a learnable multi-modal feature fusion framework, called AutoAlign, which enhances the fusion process at both pixel level and instance level.
- We present a joint training paradigm for 2D-3D detection to regularize the features extracted from the image branch and improve the detection accuracy.
- Through extensive experiments, we validate the effectiveness of the proposed AutoAlign on various 3D detectors and achieve competitive performance on both KITTI and nuScenes datasets.

2 Related Work

2.1 3D Object Detection with Single Modality

3D object detection is often conducted through a single modality of either RGB camera or LiDAR sensor. Camera-based 3D methods take the image as input and output the localization of objects in the space. Since monocular cameras cannot provide depth information, these models need to estimate the depth themselves [Chen *et al.*, 2016]. For example, [Mousavian *et al.*, 2017] first predicts the 2D bounding boxes and then estimates the depth of objects to unfold 2D boxes into 3D. However, monocular 3D detection often fails at predicting depth information. Therefore, stereo images are utilized to generate dense point clouds for 3D de-

tection [You *et al.*, 2019; Li *et al.*, 2019]. The most widely used sensors for 3D detection are LIDARs, which can be categorized into three categories: voxel, point, and view. Voxel-based techniques discretize points into voxels and aggregate points into them to extract features [Zhou and Tuzel, 2018]. Different from voxel-based approaches, [Shi *et al.*, 2019; Yin *et al.*, 2021a] directly process features at the point level, which maintains the original geometrical information provided by raw points, but they are generally computationally expensive. Extracting features from each view is also a popular stream in 3D detection, where points are compressed into bird-eye view [Lang *et al.*, 2019] or range view [Fan *et al.*, 2021] for instance prediction.

2.2 3D Object Detection with Multi-modalities

Recently, multi-modal fusion for object detection attracts numerous attentions. For example, [Qi *et al.*, 2018] predicts boxes in 2D domain and further refines them in the 3D space. [Ku *et al.*, 2018] and [Chen *et al.*, 2017] attempt to perform RoI-wise fusion. In order to get more smooth BEV maps, [Yoo *et al.*, 2020] proposes to learn an auto-calibrated projection for different modalities. However, it suffers from the problem of feature blurring. Other methods [Sindagi *et al.*, 2019; Liang *et al.*, 2018] fuse features in a point-wise manner. For instance, [Vora *et al.*, 2020] paints the 2D semantic predictions on 3D points using camera projection matrix and then performs 3D object detection. [Huang *et al.*, 2020] designs a novel L_1 -Fusion module for fine-grained fusion.

3 Method

In this section, we describe the proposed AutoAlign in detail. An overview of our approach is presented in Figure 2.

3.1 Pixel-level Feature Aggregation

Previous work mainly utilizes the camera projection matrix to align image and point features in a deterministic manner. This kind of approaches are effective but may bring in two potential problems: 1) the point cannot get a broader view on image data and 2) only positional consistency is maintained while ignoring semantic correlation. Therefore, we devise the Cross-Attention Feature Alignment (CAFA) module to adaptively align the features between non-homogenous representations. Instead of adopting a one-to-one matching pattern, the CAFA module enables each voxel to perceive the whole image and dynamically attend pixel-wise 2D features based on learnable alignment maps.

As shown in Figure 2, our method uses ResNet-50 as the backbone to extract global feature maps from given images. As a result, an input image with a size of $H \times W$ will produce the feature map with the spatial dimension of $H/32 \times W/32$. The feature map extracted from the image backbone is denoted as $\mathbf{Z} \in \mathbb{R}^{h \times w \times c}$, where h, w, c are the height, width, and channel of the global feature map, respectively. An 1×1 convolution is added to reduce the feature dimension, creating a new feature map $\mathbf{F} \in \mathbb{R}^{h \times w \times d}$. After that, we flatten the spatial dimensions of F into one dimension, resulting in a $hw \times d$ feature vector. In our cross-attention mechanism, given the feature map $F = \{f_1, f_2, \dots, f_{hw}\}$ (f_i indicates the

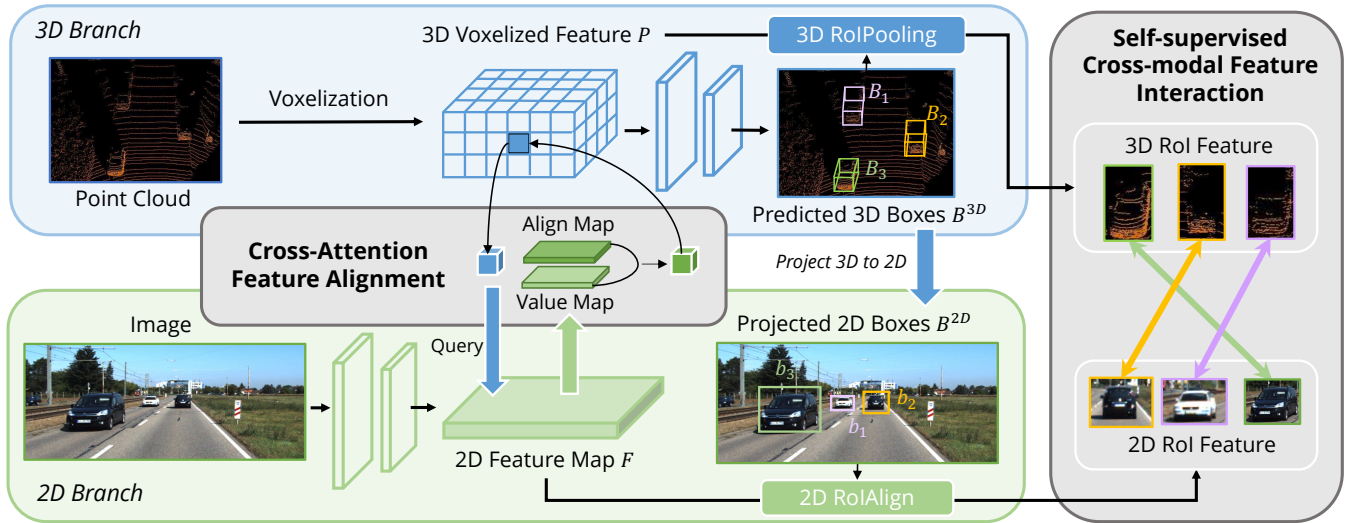


Figure 2: The framework of AutoAlign. It consists of two core components: CAFA (Sec. 3.1) performs feature aggregation on the image plane to extract fine-grained pixel-level information for each voxel feature and SCFI (Sec. 3.2) conducts cross-modal self-supervised supervision, exerting instance-level guidance to strengthen the semantic consistency in the CAFA module.

image feature of the i^{th} spatial position) and voxel features $P = \{p_1, p_2, \dots, p_J\}$ (p_j indicates each non-empty voxel feature) extracted from raw point clouds, keys and values are generated from F and queries are produced by P . Formally,

$$\mathbf{Q}_j = p_j \mathbf{W}^Q, \mathbf{K}_i = f_i \mathbf{W}^K, \mathbf{V}_i = f_i \mathbf{W}^V, \quad (1)$$

where $\mathbf{W}^Q \in \mathbb{R}^{d \times d_k}$, $\mathbf{W}^K \in \mathbb{R}^{d \times d_k}$, and $\mathbf{W}^V \in \mathbb{R}^{d \times d_v}$ are linear projections. For the j^{th} query \mathbf{Q}_j , the attention weights are calculated based on the dot-product similarity between the cross-modal query and the key:

$$s_{i,j} = \frac{\exp(\beta_{i,j})}{\sum_{j=1}^{hw} \exp(\beta_{i,j})}, \beta_{i,j} = \frac{\mathbf{Q}_j \mathbf{K}_i^T}{\sqrt{d_k}}, \quad (2)$$

where $\sqrt{d_k}$ is a scaling factor. The output of the cross-attention mechanism is defined as the weighted sum over all values according to the attention weights:

$$\hat{f}_i^{\text{att}} = \text{Att}(\mathbf{Q}_i, \mathbf{K}, \mathbf{V}) = \sum_{j=1}^{hw} s_{i,j} \mathbf{V}_j. \quad (3)$$

The normalized attention weight $s_{i,j}$ models the interests between different spatial pixels f_i and voxel p_j , which is the *align map* shown in Figure 2. The weighted sum of the values can aggregate fine-grained spatial pixels to update p_j , which enriches the point features with 2D information in a global view manner. Like the transformer architecture, we use the feed-forward network to produce the final RGB-aware point features as:

$$\mathbf{F}^{\text{att}} = \text{FFN}(\hat{\mathbf{F}}^{\text{att}}), \quad (4)$$

where $\text{FFN}(\cdot)$ is a simple neural network using one fully-connected (FC) layer [Vaswani *et al.*, 2017].

3.2 Instance-level Feature Interaction

The CAFA is a fine-grained paradigm in aggregating image features. However, it fails at capturing instance-level information. On the contrary, RoI-wise feature fusion maintains

the integrity of the object while suffers from its coarse feature aggregation and the absence of 2D information during the proposal generation phase.

To bridge the gap between pixel-level and instance-level fusions, we introduce the Self-supervised Cross-modal Feature Interaction (SCFI) module to guide the learning of the CAFA. It directly utilizes the final predictions of 3D detector as proposals, which leverages both image and point features for accurate proposal generation. Moreover, instead of concatenating cross-modal features together for further box refinement, we conduct the similarity constraint between paired cross-modal features, as a manner of instance-level guidance for feature alignment.




Figure 3: The architecture of Self-supervised Feature Interaction. Non-homogenous ROI features from images and points are both processed by MLP head to generate cross-modal representation for feature interaction.

Given 2D feature map \mathbf{F} and the corresponding 3D voxelized features \mathbf{P} , we randomly sample N regional 3D de-

tection boxes, denoted as $\mathbf{B}^{3D} = (B_1^{3D}, B_2^{3D}, \dots, B_N^{3D})$, and then project them into the 2D plane using the camera projection matrix, resulting in a set of paired 2D boxes $\mathbf{B}^{2D} = (B_1^{2D}, B_2^{2D}, \dots, B_N^{2D})$. Once obtaining the paired boxes, we adopt 2DRoIAlign [He *et al.*, 2017] and 3DRoIPooling [Shi *et al.*, 2020] in both 2D and 3D feature spaces to obtain respective RoI features \mathbf{R}^{3D} and \mathbf{R}^{2D} , where each R_i^{3D} and R_i^{2D} are given by:

$$\begin{aligned} R_i^{3D} &= \text{3DRoIPooling}(\mathbf{P}, B_i^{3D}), \\ R_i^{2D} &= \text{2DRoIAlign}(\mathbf{F}, B_i^{2D}). \end{aligned} \quad (5)$$

For each paired 2D and 3D RoI features, we perform self-supervised cross-modal feature interaction on C5 from the image branch and the feature after voxelization from the point branch. Both of them are fed into a projection head h , transforming the output of one modality to match another modality. Similar to [Chen and He, 2021], a prediction head f with two fully connected layers is introduced. Denoting the two output vectors as $p_1 = f(h(R^{3D}))$ and $q_2 = h(R^{2D})$, we minimize the feature distance $\mathcal{D}(p_1, q_2)$ with negative cosine similarity loss, as illustrated in Figure 3. To pull the two modal representations closer, we define a symmetry loss as:

$$\mathcal{L}_{SCFI} = \frac{1}{2}\mathcal{D}(p_1, q_2) + \frac{1}{2}\mathcal{D}(p_2, q_1). \quad (6)$$

Besides, the stop-gradient strategy is also adopted for the branch without prediction heads, which can be represented by $D(m_1, \text{stopgrad}(v_1))$. Hence, the interaction loss is implemented as:

$$\mathcal{L}_{SCFI} = \frac{1}{2}D(p_1, \text{stopgrad}(q_2)) + \frac{1}{2}D(p_2, \text{stopgrad}(q_1)). \quad (7)$$

3.3 Joint Training for 2D-3D Detection

Despite the effectiveness of multi-task learning, there is little work discussing the joint-detection for both image and point domains. In most previous methods, the image backbone is directly initialized with the pre-trained weights from other external datasets. During the training phase, the only supervision is the 3D detection loss, propagated from the point branch. Considering the large number of parameters in the image backbone, the 2D branch is more likely to get overfitting with implicit supervision. To regularize the representation extracted from the image, we extend the image branch into Faster R-CNN and supervise it with 2D detection loss, where the total loss \mathcal{L} is designed as:

$$\mathcal{L} = \mathcal{L}_{3D} + \mathcal{L}_{2D} + \mathcal{L}_{SCFI}, \quad (8)$$

$$\mathcal{L}_{3D} = \mathcal{L}_{3D}^{cls} + \mathcal{L}_{3D}^{reg}, \quad (9)$$

$$\mathcal{L}_{2D} = \mathcal{L}_{rpn}^{cls} + \mathcal{L}_{rpn}^{reg} + \mathcal{L}_{rcnn}^{cls} + \mathcal{L}_{rcnn}^{reg}. \quad (10)$$

4 Experiments

4.1 Implementation Details

To validate the effectiveness of our AutoAlign, we select PointPillar [Lang *et al.*, 2019], SECOND [Yan *et al.*, 2018],

and CenterPoint [Yin *et al.*, 2021a] as representative methods for our experiments. For the image branch, Faster R-CNN [Ren *et al.*, 2015] with ResNet50 is adopted as the 2D detector. The hidden units of cross-attention alignment module are set to 128 and the output sizes of 2DRoIAlign and 3DRoIPooling are both set to 4. The MLP units of the projector and predictor of the self-supervised cross-modal module are 2048 and the hidden unit number is 512. Our 2D-3D joint training framework is optimized in an end-to-end manner with hybrid optimizers where the 3D branch is optimized with AdamW and the 2D branch is optimized with SGD. We use MMDetection3D [Contributors, 2020] as our codebase, and apply the default settings if not specified.

4.2 Results on KITTI dataset

In this section, we evaluate our framework on the KITTI dataset and report the average precision (AP₄₀). We implement AutoAlign on two representative 3D object detectors: PointPillar (pillar-based) and SECOND (voxel-based). The 3D mAP performance is reported in Table 1. Overall, our AutoAlign significantly improves PointPillar and SECOND by 3.0 and 2.3 mAP under 3D moderate evaluation protocol, which validates the effectiveness of the proposed method. When observing the results in detail, we find that the APs of Pedestrian and Cyclist are promoted most (3.0 and 3.5 mAP on AP_{3D} moderate, respectively). We infer the reason that cars often hold more points, while objects like pedestrians and cyclists are mostly short of reflection, which makes them harder to be detected in the 3D space. Therefore, AutoAlign benefits from the RGB data that are naturally dense and rich in semantic and texture information.

4.3 Results on NuScenes dataset

We also conduct experiments on the much larger nuScenes dataset with current state-of-the-art 3D detector CenterPoint to further validate the effectiveness of AutoAlign. As shown in Table 3, AutoAlign achieves 66.6 mAP and 71.1 NDS on the nuScenes validation set, outperforming the strong CenterPoint baseline by 7.0 mAP and 6.5 NDS. It also surpasses the recently developed multi-modal 3D detector MVP [Yin *et al.*, 2021b] by 1.1 NDS under the same single-stage settings. Besides, due to its simplicity and joint-training paradigm, it does not require any sophisticated virtual point generation or image feature pre-fetching, which is much more suitable for real-world applications. We also report detailed results on each object category as well as the performance on the test leaderboard in the supplementary material.

4.4 Ablation Studies

To understand how each module in AutoAlign promotes the detection accuracy, we test each component on the baseline detector SECOND and report its AP performance on the KITTI validation dataset in Table 2.

When cross-attention feature alignment is applied, the accuracy is raised by 0.5 mAP and the improvements are found on objects of all difficulty levels. This result validates the importance of preserving the high resolution of image information when aggregating cross-modal features.

Method	AutoAlign	Car AP _{3D} (%)			Pedestrian AP _{3D} (%)			Cyclist AP _{3D} (%)			Overall AP _{3D} (%)		
		Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
PointPillar	✓	85.89	73.88	67.97	50.17	45.10	41.09	78.66	59.51	56.01	71.57	59.50	55.02
		87.13	75.48	69.87	54.87	48.53	44.61	82.25	63.40	58.89	74.75	62.47	57.79
SECOND	✓	87.80	77.47	74.68	64.73	59.08	52.84	83.56	67.42	62.97	78.70	67.99	63.50
		88.16	78.01	74.90	69.67	62.03	58.59	86.04	70.89	65.83	81.29	70.31	66.44

Table 1: AP_{3D} performance on different 3D object detectors w/o and w/ AutoAlign on KITTI validation set.

CAFA (pixel-level)	SCFI (instance-level)	2D Joint Training	AP _{3D} (%)			AP _{BEV} (%)		
			Easy	Mod.	Hard	Easy	Mod.	Hard
			78.70	67.99	63.50	81.05	74.42	70.86
✓			79.64	68.54	64.24	81.35	75.13	71.34
✓	✓		80.63	69.67	65.49	82.13	76.04	72.53
✓	✓	✓	81.29	70.31	66.44	83.68	77.71	73.72

Table 2: Effect of each component in our AutoAlign. Results are reported on KITTI validation set with SECOND.

Method	mAP	NDS	Reference
PointPainting	45.6	54.6	CVPR 2020
3D-CVF	42.1	49.8	ECCV 2020
CenterPoint	56.4	64.8	CVPR 2021a
ObjectDGCNN	58.6	66.0	NeurIPS 2021
MVP	66.0	70.0	NeurIPS 2021b
CenterPoint*	59.6	66.6	CVPR 2021a
CenterPoint + AutoAlign	66.6	71.1	-

Table 3: mAP and NDS performance on nuScenes dataset. The models are trained on nuScenes train subset and evaluated on nuScenes validation subset. * indicates our re-implementation.

Then, we add the SCFI module which brings in a 1.2 mAP enhancement, namely the overall moderate AP_{3D} improves from 68.5 to 69.7, suggesting that feature interaction plays a pivotal role in our fusion framework. It exerts instance-level supervision on the feature alignment, which hints at how to aggregate semantically paired features across non-homogenous representations.

When 2D joint training is added, the accuracy is boosted by another 0.6 mAP and the AP_{hard} is raised by 1.0 mAP. Such a large improvement benefits from two aspects: 1) the joint training paradigm regularizes the optimization of image backbone and 2) joint optimization reduces the training gap between 2D and 3D models and maintains the feature consistency during the cross-modal feature fusion process.

4.5 Discussions

In this section, we delve into AutoAlign framework to study how the detection accuracy is achieved and gain a deeper understanding of the underlying mechanisms. For all experiments, we employ SECOND with the same settings in Section 4.2.

Investigating the Best Cross-modal Query Strategy.

In this part, we compare various strategies for cross-modal feature query. Firstly, we choose the widely adopted fusion strategy, *i.e.*, projecting points to the image plane through the

camera projection matrix and utilizing point-wise bilinear-interpolation to obtain the aligned 2D image. Secondly, we test the non-local block proposed in [Wang *et al.*, 2018], where all image features are taken into account but only interested positions with high attention scores will be aggregated for cross-modal fusion. Finally, we adopt a more general form, which is similar to the self-attention module in [Vaswani *et al.*, 2017], but we extend it from the same modalities to non-homogenous representations. Following the common design of self-attention, we explore the performance of single-head cross-attention module and the multi-head one. The detailed results are listed in Table 4. When using point-based projection, the improvement is limited, since the points are unable to get continuous image features. However, when replacing the point-based projection with non-local block, the performance is still unsatisfying. The possible reason may lie in the FC layers that lead to the overfitting issue. Compared to vanilla non-local block, the performance of cross-attention is more competitive, probably due to the adoption of dropout strategy and feature normalization. Considering the computational cost and efficiency, we finally take single-head cross-attention as our query strategy.

Query Strategy	AP _{3D} (%)		
	Easy	Mod.	Hard
Point-based Proj	80.24	69.40	65.65
Non-Local	80.01	69.13	65.34
Multi-head Cross-attention	81.04	70.25	66.49
Single-head Cross-attention	81.29	70.31	66.44

Table 4: AP_{3D} performance with various query strategies for cross-modal feature alignment.

Seeking the Suitable Feature Source for Self-supervised Feature Interaction.

Feature interaction is a core component since it intensifies the semantic consistency of the CAFA module with instance-level guidance. Hence, how to select suitable feature source for self-supervised learning is non-trivial. After carefully examining the selection of point feature and image feature




Figure 4: Visualization of alignment maps generated by CAFA module from two randomly selected point voxels. To validate the effectiveness of SCFI module, we also visualize the alignment map without SCFI module. SCFI regularizes CAFA with instance-level semantic supervision, resulting in a positionally and semantically meaningful alignment map.

sources, we take the image features directly from ResNet backbone (*i.e.*, C5) and after FPN (*i.e.*, P5) as candidates. For the point branch, we select the features before point backbone and after backbone. As shown in Table 5, using C5 as the image feature is better than P5. We infer the reason that P5 is directly for 2D detection and therefore limits the generalization ability for cross-modal feature fusion, while C5 is more flexible for both 2D detection and non-homogenous self-supervised learning. When choosing point features after backbone, we observe a quick convergence of the similarity loss, but the result is unsatisfying. This may stem from the too much flexibility of the 3D branch, *i.e.*, the point backbone provides possibility of complex transformation for point features, which eases the optimization of loss but weakens the instance-level guidance of semantic consistency by our proposed self-supervised feature interaction. On the contrary, although using the features before backbone slows down the convergence, the model is implicitly supervised by mutual interaction and gradually learns how to align cross-modal features in the CAFA module.

Img Feat	Pts Feat	AP _{3D} (%)		
		Easy	Mod.	Hard
P5	after backbone	78.42	67.49	63.98
C5	after backbone	79.14	68.28	64.73
P5	before backbone	81.53	70.06	66.11
C5	before backbone	81.29	70.31	66.44

Table 5: AP_{3D} performance with different feature sources from image/points for cross-modal feature interaction.

Optimal Loss for Self-supervised Cross-modal Learning.

Since most self-supervised learning approaches are based on homogenous representations, exploring the optimal self-supervised loss for cross-modality is necessary. We compare four different prototypes and report the results in Table 6. We adopt the classical version of contrastive loss, where both positive and negative pairs are considered. Note that features located at the same position in 3D space and 2D plane are considered as positive pairs while the rest ones are negative pairs. The selections of NCE loss and its variant InfoNCE do not provide remarkable enhancement. However, when utilizing positive pairs only for feature interaction, we observe significant improvements. We infer the reason that the points

hold less identity information compared to images. When supervising two similar instances with negative pair loss, it may deteriorate the 3D feature representation if the shapes of the instances are similar to each other. Therefore, we choose the negative cosine similarity loss with positive pairs for our feature interaction module.

Loss	Pair	AP _{3D} (%)		
		Easy	Mod.	Hard
NCE Loss	Pos + Neg	79.14	68.65	65.54
InfoNCE Loss	Pos + Neg	79.63	68.57	65.21
CE Loss	Pos	80.83	69.73	66.01
NCS Loss	Pos	81.29	70.31	66.44

Table 6: AP_{3D} performance with different loss settings for self-supervised cross-modal learning. (NCS: negative cosine similarity)

4.6 Visualization and Analysis

In addition to the detection results on various datasets, we also provide a more straightforward visualization on the learned alignment map, given different 3D query voxel features, as shown in Figure 4. To better illustrate the mutual influence between the CAFA and SCFI modules, we compare the query attention map with and without SCFI. It can be clearly concluded that CAFA fails to yield a meaningful alignment map on 2D images without feature interaction. Conversely, when armed with SCFI, the CAFA module can successfully provide a positionally and semantically reasonable feature alignment map.

5 Conclusion

In this work, we develop AutoAlign, a learnable multi-modal feature fusion method for 3D object detection. The proposed Cross-Attention Feature Alignment module enables each voxel feature to aggregate image information in a fine-grained manner. Furthermore, a novel Self-supervised Cross-modal Feature Interaction module is designed to enhance the semantic consistency during assignment for the CAFA module. Comprehensive experimental results have demonstrated that AutoAlign significantly improves various 3D detectors on the KITTI and nuScenes datasets. We hope our work could provide a new perspective in multi-modal feature fusion for autonomous driving.

References

- [Chen and He, 2021] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *CVPR*, pages 15750–15758, 2021.
- [Chen *et al.*, 2016] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *CVPR*, pages 2147–2156, 2016.
- [Chen *et al.*, 2017] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *CVPR*, pages 1907–1915, 2017.
- [Contributors, 2020] MMDet3d Contributors. MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>, 2020.
- [Fan *et al.*, 2021] Lue Fan, Xuan Xiong, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Rangedet: In defense of range view for lidar-based 3d object detection. *arXiv preprint arXiv:2103.10039*, 2021.
- [He *et al.*, 2017] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, pages 2961–2969, 2017.
- [Huang *et al.*, 2020] Tengpeng Huang, Zhe Liu, Xiwu Chen, and Xiang Bai. Epnet: Enhancing point features with image semantics for 3d object detection. In *ECCV*, pages 35–52, 2020.
- [Ku *et al.*, 2018] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *IROS*, pages 1–8, 2018.
- [Lang *et al.*, 2019] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, pages 12697–12705, 2019.
- [Li *et al.*, 2019] Peiliang Li, Xiaozhi Chen, and Shaojie Shen. Stereo r-cnn based 3d object detection for autonomous driving. In *CVPR*, pages 7644–7652, 2019.
- [Liang *et al.*, 2018] Ming Liang, Bin Yang, Shenlong Wang, and Raquel Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *ECCV*, pages 641–656, 2018.
- [Mousavian *et al.*, 2017] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *CVPR*, pages 7074–7082, 2017.
- [Pang *et al.*, 2020] Su Pang, Daniel Morris, and Hayder Radha. Clocs: Camera-lidar object candidates fusion for 3d object detection. In *IROS*, pages 1–10, 2020.
- [Qi *et al.*, 2018] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, pages 918–927, 2018.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *NeurIPS*, 28:91–99, 2015.
- [Shi *et al.*, 2019] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, pages 770–779, 2019.
- [Shi *et al.*, 2020] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *TPAMI*, 2020.
- [Sindagi *et al.*, 2019] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvx-net: Multimodal voxelnet for 3d object detection. In *ICRA*, pages 7276–7282, 2019.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 5998–6008, 2017.
- [Vora *et al.*, 2020] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *CVPR*, pages 4604–4612, 2020.
- [Wang and Solomon, 2021] Yue Wang and Justin M Solomon. Object dgcnn: 3d object detection using dynamic graphs. *NeurIPS*, 34, 2021.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, pages 7794–7803, 2018.
- [Xie *et al.*, 2020] Liang Xie, Chao Xiang, Zhengxu Yu, Guodong Xu, Zheng Yang, Deng Cai, and Xiaofei He. Pi-r-cnn: An efficient multi-sensor 3d object detector with point-based attentive cont-conv fusion module. In *AAAI*, volume 34, pages 12460–12467, 2020.
- [Yan *et al.*, 2018] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, pages 3337–2247, 2018.
- [Yin *et al.*, 2021a] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *CVPR*, pages 11784–11793, 2021.
- [Yin *et al.*, 2021b] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multimodal virtual point 3d detection. *NeurIPS*, 34, 2021.
- [Yoo *et al.*, 2020] Jin Hyeok Yoo, Yecheol Kim, Jisong Kim, and Jun Won Choi. 3d-cvf: Generating joint camera and lidar features using cross-view spatial feature fusion for 3d object detection. In *ECCV*, pages 720–736, 2020.
- [You *et al.*, 2019] Yurong You, Yan Wang, Wei-Lun Chao, Divyansh Garg, Geoff Pleiss, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar++: Accurate depth for 3d object detection in autonomous driving. *arXiv preprint arXiv:1906.06310*, 2019.
- [Zhou and Tuzel, 2018] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, pages 4490–4499, 2018.