





Learned Fusion: 3D Object Detection Using Calibration-Free Transformer Feature Fusion

Michael Fürst^{1,2}^a, Rahul Jakkamsetty²^b, René Schuster^{1,2}^c and Didier Stricker^{1,2}^d

¹Augmented Vision, RPTU, University of Kaiserslautern-Landau, Kaiserslautern, Germany

²DFKI, German Research Center for Artificial Intelligence, Kaiserslautern, Germany
{firstname.lastname}@dfki.de

Keywords: 3D Object Detection, Calibration Free, Sensor Fusion, Transformer, Self-Attention.

Abstract: The state of the art in 3D object detection using sensor fusion heavily relies on calibration quality, difficult to maintain in large scale deployment outside a lab environment. We present the first calibration-free approach for 3D object detection. Thus, eliminating complex and costly calibration procedures. Our approach uses transformers to map features between multiple views of different sensors at multiple abstraction levels. In an extensive evaluation for object detection, we show that our approach outperforms single modal setups by 14.1% in BEV mAP, and that the transformer learns mapping features. By showing calibration is not necessary for sensor fusion, we hope to motivate other researchers following the direction of calibration-free fusion. Additionally, resulting approaches have a substantial resilience against rotation and translation changes.

1 INTRODUCTION

Environment perception is one of the pillars of advances in automated driving. Specifically, 3D object detection is critical, as knowing the position of objects in the world relative to the ego vehicle is needed for path planning and avoiding collisions.

Many object detectors apply sensor fusion to increase the average precision over single modal approaches. For example the lidar samples a far away object only with as few as 1-7 points, rendering it hardly detectable from lidar. In the camera image, the same object typically spans an area of more than 30×30 pixels and therefore can be recognized.

The current state of the art in 3D object detection uses calibration in the form of a transform- and projection-matrix. Since features are projected from one view to the other, the calibration needs to be very precise. For example an angular error of 1° results in a misalignment of 0.7m at a distance of 40 meters, the typical size of a pedestrians bounding box.

Whilst in benchmark datasets high quality calibration is given, it is difficult to obtain and maintain high quality calibration at a production scale. Since there






Figure 1: Attention for highlighted grid cells (left) are overlaid over the BEV lidar (right). The cone corresponding to the grid cell has a high attention (yellow), while the rest has low attention (black). Our calibration-free approach learns the correspondence of image and bird's-eye-view (BEV).


is variance in production, calibration must be done per car. Furthermore, during the lifetime of a vehicle deformations and thus changes in the calibration can happen due to heat, vibration and even replacement of sensors or defective parts. For calibration, typically special environments with markers are required. If a vehicle has to be regularly re-calibrated in a special environment, this poses a substantial challenge to automated driving at scale.


Instead of improving calibration or introducing continuous calibration adding complexity, we see the solution in eliminating the need for calibration. Thus, we propose and contribute:

- The new category of calibration-free sensor fusion for object detection,

^a <https://orcid.org/0000-0001-6647-5031>

^b <https://orcid.org/0009-0000-0711-229X>

^c <https://orcid.org/0000-0001-7055-9254>

^d <https://orcid.org/0000-0002-5708-6023>

- a concrete implementation using transformers, exploiting the characteristics of self-attention,
- an analysis of the effectiveness, showing that the fusion can actually be learned (see Figure 1).

2 RELATED WORK

Object detection for autonomous vehicles is a very actively researched field with many approaches. Thus, we focus on the approaches, which we consider the most influential and closely related to our work. We further categorize the work in two groups: 3D object detectors and transformers in detection.

2.1 3D Object Detectors

Monocular (RGB-only) 3D detection is popular due to the cheap price of cameras. Approaches like MergeBox (Gähler et al., 2018), (Mousavian et al., 2017) and Direct 3D Detection (Weber et al., 2019) are some of the first to push the performance envelope for monocular detection. These early approaches have difficulties with the depth perception due to the inherent depth ambiguity of monocular data.

The current state of the art research like (Zong et al., 2023), (Wang et al., 2023) and (Liu et al., 2023a) focuses on temporal information and its efficient use for detection to partially overcome depth ambiguity. However, their prediction quality still lacks behind approaches using lidar.

PseudoLiDAR (Wang et al., 2019) discovered, that convolution of depth information in the camera view is sub-optimal. Leading to us using BEV features as the primary source for BEV object detection.

LiDAR-based detection achieves much higher mean average precision (mAP) than monocular approaches. PointPillars (Lang et al., 2019) is one of the first approaches for automated driving to combine point cloud processing with efficient CNN backbones.

CenterPoint (Yin et al., 2021) successfully extends the concept of CenterNet (Duan et al., 2019) to 3D object detection. At the core, they predict the center point of the box and the size and orientation. For the centerpoint, a heat map per object class is used and allows for easy and precise localization of objects.

The current state of the art researches temporal aspects (Koh et al., 2023), scene synthesis (Zhan et al., 2023), focusing on hard samples (Chen et al., 2023) and innovative kernels (Chen et al., 2023). Techniques used here might be applicable in learned fusion as well in future research.

Sensor Fusion for object detection uses lidar and RGB for better performance and robustness by lever-

aging the strengths of each sensor. Image data is strong at initial recognition of objects, while lidar is strong at precise localization of objects in 3D.

F-PointNet (Qi et al., 2018) first uses a 2D detector to then crop a frustum in the point cloud using the calibration and predict the 3D object in that frustums point cloud. PointPainting (Vora et al., 2020) predicts segmentation masks using the camera and then colorizes the point cloud to predict the boxes there.

AVOD (Ku et al., 2018) projects anchors and proposals into the views using the calibration to allow them to crop and concatenate the features. The fused features are used for precise bounding box prediction. LRPD (Fürst et al., 2020) uses instance segmentation to generate proposals and projects them to the views, fusing similar to AVOD. It leverages RGB for initial recognition and both sensors for the fine localization.

Following the idea of explicit handling of sensors, BEV Fusion (Liu et al., 2023b) transforms camera features efficiently to the bird's-eye-view and stacks them with lidar features to predict bounding boxes. The mapping is pre-computed using the calibration.

All current sensor fusion approaches in 3D object detection use calibration matrices. The most common use is to map information from one view to another.

2.2 Transformers in Detection

Transformers introduced by (Vaswani et al., 2017) use self-attention. The self-attention is a dot product of the query and key vectors creating the attention matrix. The attention matrix is then multiplied with the values resulting in a weighted sum of the values.

Detection Transformer (DeTr) (Carion et al., 2020) uses the features of a convolutional encoder as inputs to a transformer encoder-decoder architecture. In contrast to common practice in detection, DeTr predicts a set of bounding boxes instead of boxes for each pixel or grid cell in the image. This makes the approach very general, but more difficult to train.

Similar to DeTr, TransFusion (Bai et al., 2022) encodes features using regular backbones and then uses queries to decode the bounding boxes. In a first step, the queries are generated from a heat map for the lidar and in a second step, by projecting the center of a query to 2D using spatially modulated cross attention.

Cross-Modal-Transformer (CMT) (Yan et al., 2023) is one of the best published approaches on nuScenes (Caesar et al., 2020). Similar to DeTr, it first extracts the features from camera and lidar using encoders. Then it concatenates the features and applies a transformer decoder. The queries are computed from 3D points, projected to the respective views. Thus enabling the correlation with the position embedding of

the respective views. We found in preliminary experiments, that the model does not converge if the calibration is omitted from this step.

Finally, TransFuser (Chitta et al., 2023) uses transformers to fuse features for learning navigation in a simulated environment. The TransFusion module introduced by them is very flexible and does not use calibration. However, their overall architecture produces global features unsuited for object detection.

Object detectors rely on calibration to map information between views and TransFuser cannot be applied to detection without substantial modification.

3 APPROACH

Current state of the art detectors using sensor fusion rely on calibration, even when using transformers. It is very difficult to successfully train a completely calibration-free fusion approach, since "object queries might attend to visual regions unrelated to the bounding box to be predicted, leading to a long training time for the network" (Bai et al., 2022). Indeed, naively removing calibration from existing transformer based approaches does not work in our preliminary experiments. Hence, we focus on a simple model enabling a stable training in a timely manner to show the possibility of training such calibration-free models.

Since training transformers to map between the different views is already a hard problem, the rest of the model is kept small and straightforward. We focus on using transformers to correlate features between the views and eliminating the calibration.

3.1 TransFuseDet

Our model consists of a fusion encoder, upsampling and task-specific heads. We use two ResNets (He et al., 2016) to encode lidar and camera respectively. Transformers fuse the features at different abstraction levels, similar to TransFuser (Chitta et al., 2023). However, unlike their approach, we keep the features of BEV and camera view separated to keep the spatial interpretation of the feature maps intact. This allows us to do object detection.

The features are then upsampled to increase resolution and finally used in two CenterNet-Style decoders for BEV object detection on the BEV features and 2D object detection on the features in camera space. The main goal is BEV object detection, but the 2D object detection can be viewed as an auxiliary task to improve BEV performance, as we will show in an ablation study.

Figure 2 gives an overview over our architecture and the following sections explain more details on how the calibration-free fusion works.

3.2 Feature-Extraction: Multi-View Fusion via Transformers

To extract the features, we partially follow the methodology of TransFuser (Chitta et al., 2023). We first apply a convolutional block with pooling from the ResNets to each modality. Then, the features are fused by a transformer. However, in order to apply a transformer, the features need to be pooled, flattened and then concatenated from the different views of the sensors. After the transformer, the features must be split into the views, reshaped and upsampled, to be added to the features of the respective views.

Deviating from TransFuser, the fusion step is optional after each ResNet block. In the ablation study, we experimentally derive an optimal configuration.

Additionally, we introduce upsampling to increase the resolution again, as the native output resolution of the fusion encoder is too low for precise object detection. For upsampling we simply upscale the features, concatenate them with the higher resolution features and apply 1x1 convolutions.

3.2.1 Property of Transformers Allowing Calibration-Free Fusion

The critical component of a transformer for fusion is the attention, as it allows to correlate features independent of their spatial location. For example, in our bird's-eye-view (BEV) representation the ego-vehicle is on the left facing towards the right as shown by the little arrow in Figure 2. Thus, if we have a car which is far away in the scene, it might appear centered in the camera view, but in the lidar bird's-eye-view (BEV) it is to the far right. Typically, calibration data would be used to compute the corresponding positions in BEV and camera view and then the features would be gathered or projected to the other view. This is not done in calibration-free fusion. Here, we leverage the property of the dot product in attention to correlate similar features independent of their position.

Scaled Dot-Product Attention (Vaswani et al., 2017) consists of a dot product between the query Q and the key K , scaled using the dimensionality d_k of the keys, and then weighting the values V by the resulting (softmaxed) matrix.

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{Q \cdot K^T}{\sqrt{d_k}}\right) \cdot V \quad (1)$$

The dot product computes the similarity of two vectors. For example, our convolutional encoder finds a

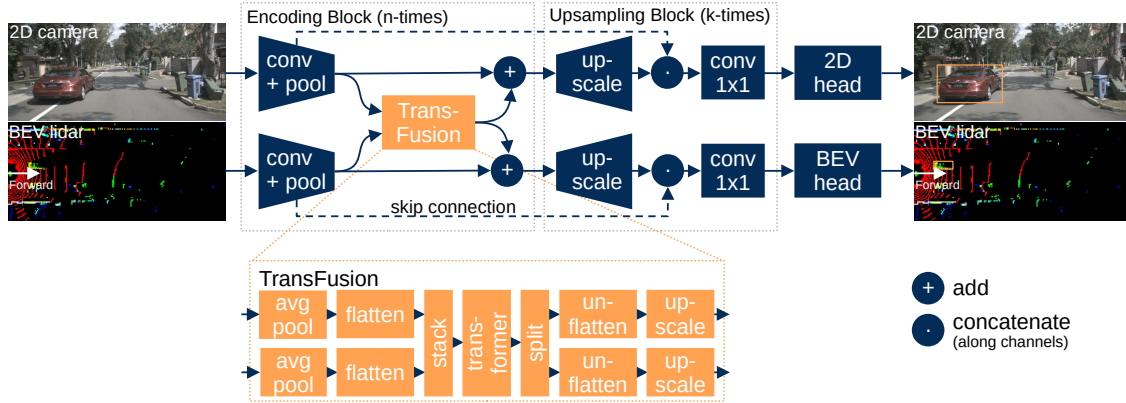


Figure 2: Our model is composed of three main components: First, the encoding block using TransFusion to fuse the features without calibration. The block is repeated $n = 4$ times and the fusion (orange) is optional in each block. Second, the upsampling block increasing the resolution of the feature maps. Upsampling is also repeated $k = 3$ times each time increasing the resolution by a factor of two. Last, a task-specific head predicting the bounding boxes in BEV and 2D.

pedestrian in the image at position x_I , and in the lidar at position x_L . Then, the resulting features f_{x_I} and f_{x_L} should have a similar semantic meaning, resulting in a high value for the dot product and thus a high attention to the feature of the other view.

3.2.2 Pooling in Fusion Reducing Computational Complexity

Transformers have a high computational cost in image processing due to the high resolution. To apply a transformer on a feature grid, the features must be flattened. Thus, a vector with shape $H \times W \times C$ must be flattened to a vector of shape $(HW) \times C$ to be used as $Q = K = V$ in self-attention. The resulting attention matrix is of shape $(HW) \times (HW) \times C$, having H^2W^2C values. So for example doubling the input resolution squares the flattened vectors size and quadratically increases the size of the attention matrix, severely limiting the resolution that is practical with respect to computational cost and learnable parameters.

For fusion, we concatenate two or more flattened feature vectors before using them as Q , K and V in the self-attention. Despite only linearly increasing complexity combined with the quadratic scaling this further increases the computational cost.

To limit cost there are two options: Reducing the resolution of the input to the model or pooling the feature maps before using them in the fusion. Pooling feature maps allows to keep high resolution feature maps and only sacrifice resolution in the fusion. For object detection with CenterNet, a higher resolution is beneficial. Higher resolution allows to distinguish objects near each other, since only a single object per grid cell can be detected. Additionally, the location of the maxima in the heat map can correspond better to the true mode, since quantization errors are smaller.

3.2.3 Upsampling

Since resolution has a big impact on the performance, we use upsampling and combine the upsampled features with the higher resolution and lower abstraction of earlier layers of the model. Following common convention, features are upsampled by a factor of two and combined with the features of the previous encoding block. The features are concatenated along the channel dimension, followed by a 1×1 convolution.

3.3 Detection Heads

The detector heads follow the proven methodology of CenterNet (Duan et al., 2019), suitable for the feature maps created by our calibration free fusion. The heads predict a class confidence and a bounding box for each grid cell in the feature maps. Additionally, in BEV we predict a yaw angle Θ .

Due to the jump in the angle from 360° to 0° , directly regressing it is not possible, as the derivative would not be well defined at this jump. Thus, we follow the convention of encoding the yaw angle Θ as a class and an offset. Specifically, we split the angle into 8 classes representing equally sized slices of the value range from -22.5° to 337.5° . This leads to the centers of each class being at 0° , 45° , 90° , ... 315° . The offset value range is then from -22.5° to $+22.5^\circ$.

To train the heads, we use a mean squared error (MSE) over the entire grid map for the heat map loss L_{heat} . Since the heat map is heavily biased towards no objects (background), we split the loss computation into background and foreground loss. These are then summed using a weighted sum:

$$L_{\text{heat}} = w_{fg}L_{\text{mse},fg} + w_{bg}L_{\text{mse},bg}. \quad (2)$$

For the box and yaw we only apply the losses on the grid cells associated with a bounding box. For other grid cells, no loss is computed. The bounding box loss L_{bbox} is a smooth L1-loss and for the yaw we use cross entropy for the class loss $L_{\Theta\text{-cls}}$ and smooth L1-loss for the offset regression loss $L_{\Delta\Theta}$.

Finally, for the BEV detection task the losses are combined into a single loss L using a weighted sum, where w denote the weights:

$$L_{\text{bev}} = w_{\text{heat}}L_{\text{heat}} + w_{\text{bbox}}L_{\text{bbox}} + w_{\Theta}(L_{\Theta\text{-cls}} + L_{\Delta\Theta}). \quad (3)$$

In the case of the 2D detection task, we use the same loss without the yaw:

$$L_{2D} = w_{\text{heat}}L_{\text{heat}} + w_{\text{bbox}}L_{\text{bbox}}. \quad (4)$$

The losses are then combined for the backpropagation using a weighted sum of BEV and 2D task loss:

$$L = w_{\text{bev}}L_{\text{bev}} + w_{2D}L_{2D}. \quad (5)$$

3.4 Implementation Details

Based on experimental studies, our model is trained with an AdamW optimizer a learning rate of $1e-4$ and weight decay of $1e-4$. The learning rate is reduced using exponential decay of 0.99996 every two steps. For the loss weights we use $w_{\text{BEV}} = 0.95$, $w_{2D} = 0.05$, $w_{fg} = 0.9$, $w_{bg} = 0.1$, $w_{\text{heat}} = 10$, $w_{\text{bbox}} = 1.0$ and $w_{\Theta} = 0.2$. The model performs best without random rotations and offsets of the lidar, but horizontal mirroring is enabled for both lidar and camera.

We do not use extensive point cloud augmentation or aggregation of points from multiple frames. This reduces our mAP significantly, but allows us to clearly show what our approach contributes instead of the influence of augmentation strategies.

4 EVALUATION

We evaluate our approach on the nuScenes dataset (Caesar et al., 2020) since it provides the required modalities, i.e. lidar and camera, while having a sufficient size. NuScenes has 40,000 annotated key frames in 1000 driving sequences, which is sufficiently large for training a transformer-based calibration-free model. However, we noticed during preliminary experiments that larger transformer-based calibration-free models have difficulties converging. Thus, we kept our model small to show the possibility of calibration-free fusion.

The frames provided by nuScenes contain 6 camera images, 1 lidar scan and 5 radars. We only use the front and rear camera, as the side cameras have many

Table 1: Learned fusion outperforms its RGB-only or lidar-only counterpart. The models are as identical as possible to eliminate all other effects except for the learned fusion.

Method	BEV mAP	2D mAP
RGB-only	-	47.0
LiDAR-only	37.6	-
Learned Fusion [ours]	42.9	48.7

Table 2: Comparison of bird’s-eye-view mAP between calibration-free fusion and much larger calibration-based approaches, which cannot be applied without calibration.

Method	Calib	mAP
CMT (Yan et al., 2023)	Yes	70.4
BEVFusion (Liu et al., 2023b)	Yes	70.2
Learned Fusion [ours]	No	48.8

frames without objects, yielding little benefit training the model. Additionally, using two instead of six cameras reduces training time by a factor of three.

4.1 Advantage over Single Sensor

Our approach introduces a new category of approaches next to camera only, lidar-only and calibration-based fusion, we introduce calibration-free fusion. Thus, comparability to existing approaches is limited.

From our approach we can simply derive a uni modal variant by removing the fusion and the sub network for the other modality. This allows us to show the concrete benefit from calibration-free fusion over single modality. However, due to the nature of our approach, it is not possible to add calibration. Thus, a direct comparison of identical approaches showing the potential advantage of calibration over calibration-free is not possible.

Our approach is significantly outperforming its two derivatives using single modality. This shows that fusion without calibration has a clear advantage over no fusion. Table 1 shows that fusion is 14.1% better than the lidar-only approach in BEV detection and 3.6% better than the camera only approach. The smaller gap for 2D detection makes sense, since camera is generally considered sufficient for 2D detection.

However, when comparing our approach to the current best approach on nuScenes using calibration, it is evident, that there is still a significant performance gap (see Table 2). This gap is to be expected from a completely novel approach using no calibration. As there is little prior work to build on, many simplifications were done in this first approach. However, we are confident, that further research adding advanced data augmentations and training strategies

Table 3: The performance impact different fusion configurations is measured. Adding fusions after the early layers of the model does not improve performance in BEV.

Method	Fusion	BEV mAP	2D mAP
Learned Fusion	4	37.7	48.1
<i>Learned Fusion</i>	<i>3,4</i>	42.9	48.7
Learned Fusion	2,3,4	40.1	47.6
Learned Fusion	1,2,3,4	39.0	47.0

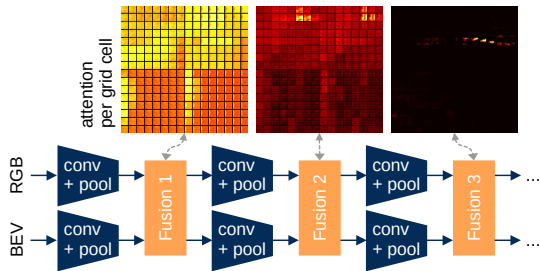


Figure 3: Visualizing the attention matrix, so that it corresponds to the grid cells in the image, shows that later fusions have more focus. Early fusions are uniformly distributed and do not contribute to a good fusion.

as well as scaling up the model can close the gap.

Without using any calibration information our approach learns to correlate the data from different views. We believe that the calibration-free nature of our approach and the advantage over single modal approaches make this a very promising new field in sensor fusion for object detection.

4.2 Number of Fusion Layers

Our approach allows to have a fusion in 4 optional places. In this study we evaluate which fusions yield the best results. We can do a fusion after each of the N last convolutional blocks in the ResNets. For example fusions at 3 & 4 means there is a fusion after the second last and last block in the ResNet of the two modalities. In Figure 3, the first three fusions of a model with all fusion blocks are shown.

In Table 3, we can see, that more fusion layers than two does not have a positive impact on the performance of the model. For example in the case of four fusions the BEV mAP degrades by 3.9 mAP ($\sim 9\%$) compared to two fusions.

A possible explanation is that the feature vectors correlated by our approach need a certain abstraction level. Lower level feature maps lack the abstraction and confuse the model. For example, a tire in an image might not be visible in the lidar at all, thus correlating these low level features does not provide much value. Adding this fusion means the model needs to actively learn to ignore the features of a different

Table 4: Random rotation and translation barely reduce the mAP of the model. Even extreme movement of the lidar by 5.5 m and rotating up to 15° only has slight impact on mAP.

Rot.	Trans.	Mirror	BEV mAP	2D mAP
0°	0m	No	41.1	47.9
0°	0m	Yes	42.9	48.7
0°	0.5m	Yes	42.3	46.7
15°	0.5m	Yes	41.3	47.1
15°	5.5m	Yes	40.5	45.3

view. Inspecting the heat maps of the attention, we noticed that the attention for the feature maps in the earlier layers is low and unfocused, see an example in Figure 3.

4.3 Sensor Displacement or Rotation

Since our model does not use explicit calibration, but correlates features, it has a built-in robustness against sensor displacement and rotation. To evaluate this, we added random rotation and translation to the input data of our model during training and testing time.

In Table 4, it is visible, that the model with mirroring performs best. However, adding random translation of up to 0.5m only reduces the mAP by 0.6 mAP in BEV. Adding random rotations of up to 15° again only reduces the mAP in BEV by 1.6. For context, calibration-based approaches expect precision in the range of 0.01m and less than 1° error. Finally, we evaluate extreme translations up to 5.5 meters (row 5) losing 2.4 mAP over the baseline (row 2). However, with an offset of 5.5 meters the sensor could be mounted on a vehicle next to ours, showing the extreme robustness to translation and rotation of the sensors in our approach.

These results show, that our approach has strong robustness against changes in the alignment of sensors with respect to translation and rotation. The correlation of features via the attention is crucial for this. We see the potential to apply this approach to multiple different vehicles without the need for modification.

4.4 Loss and Task Weighting

When training the approach there are hyperparameters regarding the weighting in the loss. We split the weighting in two categories, task weighting (see Table 5) and loss weighting (see Table 6). In task weighting, we evaluate the balance between the 2D and BEV detection loss to achieve optimal BEV mAP. For loss weighting, we evaluate the impact of the weights for heat map foreground w_{fg} , background w_{bg} and loss components w_{heat} and w_Θ .

A higher weight to the BEV loss increases per-

Table 5: Weighting the tasks differently affects performance. BEV detection is the primary task reaching best performance at 0.95 and 0.05 weighting.

w_{BEV}	w_{2D}	BEV mAP	2D mAP
0.20	0.80	38.1	45.3
0.50	0.50	41.8	49.4
0.80	0.20	42.1	49.7
0.90	0.10	40.7	47.4
0.95	0.05	42.9	48.7
0.99	0.01	41.9	44.0
1.00	0.00	40.5	0.0

Table 6: The weighting of the different loss components has a significant impact on the performance. A high heat map weight has strong impact on the BEV mAP. Our baseline (row 3) is outperformed by a very high w_{heat} (row 6).

w_{fg}	w_{bg}	w_{heat}	w_{Θ}	BEV mAP	2D mAP
0.80	0.20	10	0.2	43.5	46.8
0.90	0.10	2	0.2	39.3	45.3
0.90	0.10	10	0.2	42.9	48.7
0.90	0.10	10	1.0	44.2	50.7
0.90	0.10	20	0.2	42.8	49.4
0.90	0.10	50	0.2	48.8	49.0
0.95	0.05	10	0.2	41.5	47.3

formance (see Table 5). However, increasing beyond 0.95 reduces mAP again. The 2D loss can be viewed as an auxiliary loss for the camera feature extractor. However, reaching 40.5 mAP even without the 2D loss, the fusion model outperforms lidar only at 37.6 mAP. Fusion alone contributes substantially to the performance and the auxiliary 2D losses further improves the effectiveness of fusion.

Increasing the heat map weighting in the loss has a significant impact on the mAP (see Table 6). This can be explained by the fact, that the correct location of the center of a bounding box in the BEV is most important to detection, since the size variation within an object class is small. For example, all pedestrians have almost the same size in BEV.

4.5 Attention Map Analysis

Besides the quantitative analysis of our approach, we visualized and analyzed the behavior of the attention maps to gain understanding.

For the visualization we reformatted the attention map of the two flattened vectors into a grid of attention images for each cell in the feature map. In the Figure 3 the full matrix can be seen per layer. As described in Section 4.2 the feature maps of earlier layers are uniform and lack focus. Thus, for our subsequent analysis, we focused on the feature maps of fusion layers 3 and 4.




Figure 4: Depending on where in the image the object is, the attention visualized in BEV is different. Nearby vs far away and left vs right objects have different attention.

To validate that the attention has a meaningful interpretation, we analyzed the attention maps of grid cells containing objects. In Figure 4, we show the attention for nearby, centered, left, right and far away objects. Generally a correlation of the attention to the region in the lidar can be seen. However, in some cases the object attention lacks focus and is quite diffuse. This especially happens for nearby objects covering large areas of the image.

Analyzing many of the attention maps, we noticed a few trends: Firstly, especially the later attention layers focus their attention on cones and regions containing objects. Secondly, the attention of regions corresponding to each other have a high attention for each other. For example, the region at the horizon of the image and the far away region in the lidar have a high attention for each other.

We conclude from this observation, that the model learns correlating features from the regions of the different views containing the same objects.

4.6 Qualitative Results

To validate the plausibility of our results, we visualized the BEV detections over the pointcloud input to the model. This allows us to inspect the detections and see strengths and weaknesses of the model.

Overall, the model performance is good (see Figure 5). The model predicts the position and orientation of objects well. However, in some cases the model has false positives for pedestrian objects as well as false negatives for construction workers. The issues with construction workers are to be expected as they are an underrepresented class in the dataset.

As there is no specific shortcomings of the model apparent, our fusion seems to be quite robust.



Figure 5: Visualization of the BEV detections predicted by our model. The model has good prediction for most cars, but has difficulties at the edges of the BEV lidar. It cannot detect all construction workers in image two, however construction workers are underrepresented in the dataset. At night the detection of the cars is very good despite the low visibility.

5 CONCLUSIONS

Overall, we introduce fully calibration-free sensor fusion using neither intrinsic nor extrinsic calibration. Eliminating the need for complex calibration procedures in sensor fusion for object detection. As a calibration-free sensor fusion in object detection, we present an approach using transformers for correlating features between the views and then upscaling the features to achieve the necessary resolution for precise object detection.

In the thorough evaluation, we show that calibration-free sensor fusion is a promising field. Concretely, we show that adding calibration-free fusion increases performance over using a single modality. Further, we found that our approach is robust against random translation and rotation, since the model correlates features without a calibration matrix.

5.1 Limitations and Future Research

However, due to the complexity of the learning problem to correlate all features from the different views with each other, the model has limitations. We identified two main limitations and propose directions for

further research to eliminate them.

First, the model complexity is a problem. The transformer for the fusion requires many of parameters and thus allows for a very limited resolution. This leads to a gap in performance compared to the best calibration-based fusion approaches. We expect that approaches such as deformable-attention or a query based approach like DeTr could help here.

The second limitation is the stability of the training. During early experimentation we were experimenting with a much larger DeTr style model, but found the training to be too unstable. We attribute this to the fact that learning a correlation between random features at the beginning of the training is highly unstable. Thus, advanced strategies for training and especially pre-training could prove a very valuable direction of further research.

ACKNOWLEDGEMENTS

This work was partially funded by the Federal Ministry of Education and Research Germany under the project DECODE (01IW21001).

REFERENCES

- Bai, X., Hu, Z., Zhu, X., Huang, Q., Chen, Y., Fu, H., and Tai, C.-L. (2022). Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *IEEE/CVF conference on computer vision and pattern recognition*.
- Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., and Beijbom, O. (2020). nuscenes: A multimodal dataset for autonomous driving. In *IEEE/CVF conference on computer vision and pattern recognition*.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., and Zagoruyko, S. (2020). End-to-end object detection with transformers. In *European conference on computer vision*. Springer.
- Chen, Y., Yu, Z., Chen, Y., Lan, S., Anandkumar, A., Jia, J., and Alvarez, J. M. (2023). Focalformer3d: Focusing on hard instance for 3d object detection. In *IEEE/CVF International Conference on Computer Vision*.
- Chitta, K., Prakash, A., Jaeger, B., Yu, Z., Renz, K., and Geiger, A. (2023). Transfuser: Imitation with transformer-based sensor fusion for autonomous driving. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 45(11).
- Duan, K., Bai, S., Xie, L., Qi, H., Huang, Q., and Tian, Q. (2019). Centernet: Keypoint triplets for object detection. In *IEEE/CVF international conference on computer vision*.
- Fürst, M., Wasenmüller, O., and Stricker, D. (2020). Lrpd: Long range 3d pedestrian detection leveraging specific strengths of lidar and rgb. In *IEEE international conference on intelligent transportation systems (ITSC)*.
- Gähler, N., Mayer, M., Schneider, L., Franke, U., and Denzler, J. (2018). Mb-net: Mergeboxes for real-time 3d vehicles detection. In *IEEE Intelligent Vehicles Symposium (IV)*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*.
- Koh, J., Lee, J., Lee, Y., Kim, J., and Choi, J. W. (2023). Mgtanet: Encoding sequential lidar points using long short-term motion-guided temporal attention for 3d object detection. In *AAAI Conference on Artificial Intelligence*.
- Ku, J., Mozifian, M., Lee, J., Harakeh, A., and Waslander, S. L. (2018). Joint 3d proposal generation and object detection from view aggregation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Lang, A. H., Vora, S., Caesar, H., Zhou, L., Yang, J., and Beijbom, O. (2019). Pointpillars: Fast encoders for object detection from point clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Liu, H., Teng, Y., Lu, T., Wang, H., and Wang, L. (2023a). Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In *IEEE/CVF International Conference on Computer Vision*.
- Liu, Z., Tang, H., Amini, A., Yang, X., Mao, H., Rus, D. L., and Han, S. (2023b). Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In *IEEE International Conference on Robotics and Automation (ICRA)*.
- Mousavian, A., Anguelov, D., Flynn, J., and Kosecka, J. (2017). 3d bounding box estimation using deep learning and geometry. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Qi, C. R., Liu, W., Wu, C., Su, H., and Guibas, L. J. (2018). Frustum pointnets for 3d object detection from rgb-d data. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. (2017). Attention is all you need. *Advances in neural information processing systems*.
- Vora, S., Lang, A. H., Helou, B., and Beijbom, O. (2020). Pointpainting: Sequential fusion for 3d object detection. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Wang, S., Liu, Y., Wang, T., Li, Y., and Zhang, X. (2023). Exploring object-centric temporal modeling for efficient multi-view 3d object detection. *arXiv preprint arXiv:2303.11926*.
- Wang, Y., Chao, W.-L., Garg, D., Hariharan, B., Campbell, M., and Weinberger, K. Q. (2019). Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*.
- Weber, M., Fürst, M., and Zöllner, J. M. (2019). Direct 3d detection of vehicles in monocular images with a cnn based 3d decoder. In *IEEE Intelligent Vehicles Symposium (IV)*.
- Yan, J., Liu, Y., Sun, J., Jia, F., Li, S., Wang, T., and Zhang, X. (2023). Cross modal transformer: Towards fast and robust 3d object detection. In *IEEE/CVF International Conference on Computer Vision*.
- Yin, T., Zhou, X., and Krahenbuhl, P. (2021). Center-based 3d object detection and tracking. In *IEEE/CVF conference on computer vision and pattern recognition*.
- Zhan, J., Liu, T., Li, R., Zhang, J., Zhang, Z., and Chen, Y. (2023). Real-aug: Realistic scene synthesis for lidar augmentation in 3d object detection. *arXiv preprint arXiv:2305.12853*.
- Zong, Z., Jiang, D., Song, G., Xue, Z., Su, J., Li, H., and Liu, Y. (2023). Temporal enhanced training of multi-view 3d object detector via historical object prediction. *arXiv preprint arXiv:2304.00967*.