

Predict to Detect: Prediction-guided 3D Object Detection using Sequential Images

Sanmin Kim Youngseok Kim In-Jae Lee Dongsuk Kum
KAIST

{sanmin.kim, youngseok.kim, oliver0922, dskum}@kaist.ac.kr

Abstract

Recent camera-based 3D object detection methods have introduced sequential frames to improve the detection performance hoping that multiple frames would mitigate the large depth estimation error. Despite improved detection performance, prior works rely on naive fusion methods (e.g., concatenation) or are limited to static scenes (e.g., temporal stereo), neglecting the importance of the motion cue of objects. These approaches do not fully exploit the potential of sequential images and show limited performance improvements. To address this limitation, we propose a novel 3D object detection model, P2D (Predict to Detect), that integrates a prediction scheme into a detection framework to explicitly extract and leverage motion features. P2D predicts object information in the current frame using solely past frames to learn temporal motion features. We then introduce a novel temporal feature aggregation method that attentively exploits Bird's-Eye-View (BEV) features based on predicted object information, resulting in accurate 3D object detection. Experimental results demonstrate that P2D improves mAP and NDS by 3.0% and 3.7% compared to the sequential image-based baseline, proving that incorporating a prediction scheme can significantly improve detection accuracy.

1. Introduction

3D object detection is an essential task for building a reliable self-driving system. In recent years, camera-based 3D object detection [29, 33, 35, 38] has gained widespread attention due to the cost-effectiveness of a camera sensor and its high-resolution characteristic. However, camera-based 3D object detection still has limited performance due to the scale ambiguity caused by projecting 3D space onto a 2D image and the absence of motion cues that are difficult to capture in a single image.

Recent works have mitigated these drawbacks by leveraging multiple frames from history. Multi-frame ap-

proaches incorporate temporal information into the space domain to provide richer information. Moreover, sequence images are often readily available in real-world applications such as autonomous driving, making the use of sequence images an attractive option for performance improvements.

Previous works [12, 14, 18, 23] have used temporal images in feature-level aggregation by concatenating sequential features to merge them. On the other hand, another line of work [16, 30, 36, 41] has adopted a temporal stereo [42] to enhance depth estimation using the multi-view stereo (MVS) [9]. Although these methods have proved the effectiveness of sequence frames over a single frame, they did not thoroughly investigate into the motion cue of objects, from which the object detection would benefit by using sequence images.

Temporal images have rich motion information, which can provide critical motion features for accurate object detection. To further demonstrate the importance of motion cues in detection, we conducted experiments and evaluated the performance of the prediction-only results, which rely on motion prediction from previous frames, without using the current frame. Our findings from Table 3 indicate that the prediction-only results (P) can achieve comparable performance to the final results (P+D), reaching up to 76% and 89% in terms of mAP and NDS, respectively. The final results (P+D) denote detection results that incorporate all temporal frames, including the current frame. This experiment highlights the potential of using motion features from previous frames, which has been overlooked in prior works.

To this end, we propose a novel sequential image-based 3D object detection model that learns motion cues to improve detection accuracy. Our approach, P2D (Predict to Detect), introduces a prediction scheme into the detection task to fully exploit multi-frame image data. Specifically, P2D conducts motion prediction using previous frames to output the predicted objects' information for the current frame. In the feature aggregation module, we employ a deformable attention [46] to make a spatio-temporal feature on the basis of prediction results that contains motion features. Finally, the 3D detection head takes aggregated the

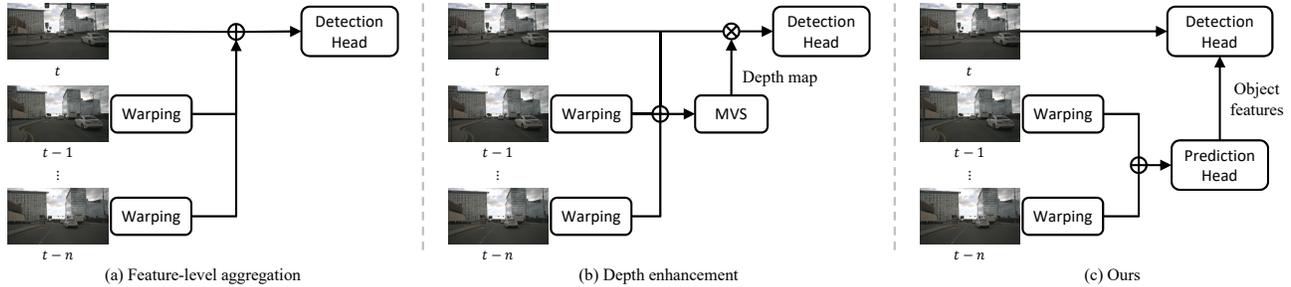


Figure 1. Comparison of temporal image-based methods. All methods align features by warping previous frames to the current frame. (a) Feature-level aggregation naively concatenate sequential features before inputting them to the detection head. (b) Depth enhancement methods facilitate depth estimation using Multi-View Stereo (MVS). (c) Our proposed approach combines prediction and detection to leverage motion features. We predict object information from previous frames and use it to detect objects in the current frame.

spatio-temporal feature and outputs the final detection results. In this way, our proposed method can fully benefit from multi-frame inputs by predicting objects’ motion and utilizing it explicitly, providing a more accurate and reliable 3D object detection system for autonomous driving.

In summary, our contributions are as follows:

- We identify the motion feature as a key factor when handling sequential images for 3D object detection. A prediction mechanism is introduced to fully exploit the motion feature of multi-frame image data.
- We propose a novel 3D object detection model using sequential images. Our model includes a Prediction Head to predict object information and a Prediction-guided Feature Aggregation to integrate temporal features using motion features.
- Our approach achieves improved performance compared to prior state-of-the-art methods. Extensive experimentation confirms the effectiveness of our approach in adapting to moving objects and accurately estimating their velocities.

2. Related Work

2.1. Camera-based 3D Object Detection

Camera-based 3D object detection has gained significant attention following the success of 2D detection methods. Early methods [2, 8, 25, 28, 29, 37, 38] exploit perspective view features by extracting 2D features from input images and directly estimating 3D information for object detection. After the pioneering work of Mono3D [6], M3D-RPN [2] proposes 3D anchor boxes and depth-aware convolution and FCOS3D [38] projects 3D targets into a 2D image plane. To mitigate the depth ambiguity of 2D images, several methods [19, 27, 32, 37] leverage geometric information while [29] employs additional depth supervision. PGD [37] constructs geometric relation graphs across predicted objects to facilitate depth estimation and GUPNet [27] estimates a depth

using height information. On the other hand, DD3D [29] boosts depth estimation ability using extra datasets [10].

Another stream of work employs view transformation to overcome the limitations of the perspective view. Several works transform image pixels into 3D point clouds using estimated depth information to take advantage of LiDAR detector [39, 44]. On the other hand, other works proposed to transform image features into voxel-like Bird’s Eye View (BEV) features for 3D perception [31, 33, 34]. BEVDet [13] uses LSS [31] based approach, which leverages the depth distribution to transform the perspective-view features into BEV space. BEVDepth [18] adds depth supervision using LiDAR point cloud, whereas BEVFormer [20] adopts a deformable attention [46].

2.2. Sequential Image-based 3D Object Detection

To improve 3D object detection performance, several works expanded the time horizon by leveraging temporally sequential frames. Sequential image-based 3D object detection can be categorized into object-centric and scene-centric methods.

Object-centric methods. Inspired by the object tracking, these methods [3, 7, 15] employ object-level association to improve detection performance. Object-centric methods detect objects frame by frame and refine the detection results by matching objects. Kinematic3D [3] uses a 3D Kalman Filter to consider the kinematic motion of objects and update detection results. MotionLoss [7] introduces patch-wise motion loss for temporal consistency. Time3D [15] adopts object-wise attention for temporal matching, where detections from the current frame operate as queries, and those from previous frames are keys and values.

Scene-centric methods. These methods are subdivided into two again: feature-level aggregation and depth enhancement. Feature-level aggregation methods [12, 14, 18, 20] extract features from each image and aggregate temporal features before inputting them to the detection head. [12, 14, 18] aggregate sequential image features by concatenating them after temporal alignment using warping into the

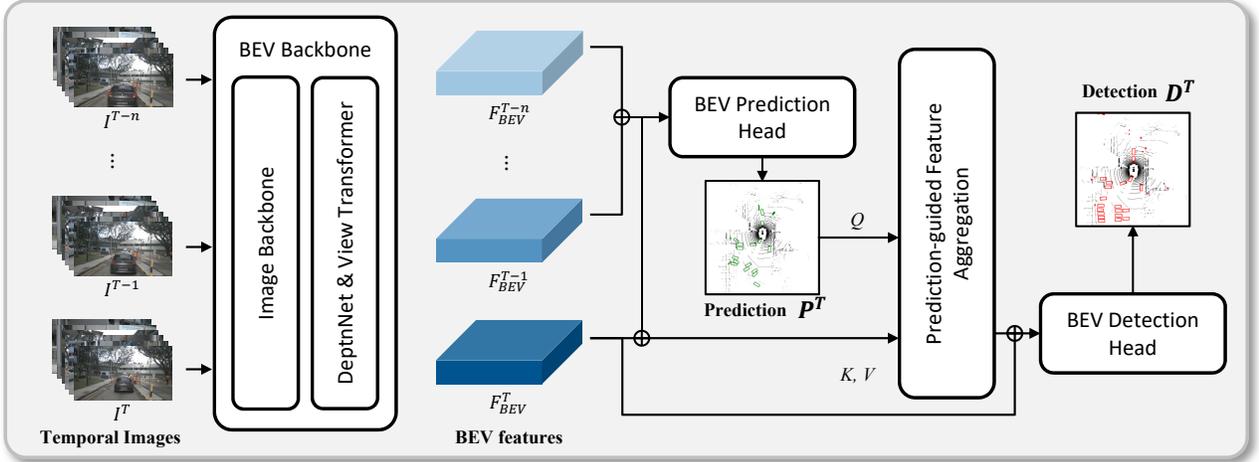


Figure 2. Overall architecture of P2D. The BEV backbone extracts BEV features from multi-view and multi-timestep images. The BEV Prediction Head takes BEV features of previous frames as the input and predicts objects in the current frame. The Prediction-guided Feature Aggregation module merges all temporal features based on predicted object information. The BEV Detection Head takes the aggregated feature and outputs the final detection results. The whole model is trained with the two loss terms of prediction and detection loss.

current timestep to compensate for the ego-motion. BEVFormer [20] adopts temporal attention to aggregate temporal features based on BEV queries. Depth enhancement methods [16, 30, 36, 41] employ temporal stereo, which extends Multi-view Stereo (MVS) [1, 42] into temporal images. By setting the translation of an ego agent as a baseline, two temporally nearby images have stereo correspondence that can be used in stereo matching. DfM [36] employs the temporal stereo in monocular 3D object detection with a theoretical analysis. BEVStereo [16] improves the temporal stereo with a sparse cost volume and an iterative algorithm inspired by MaGNet [1]. STS [41] focuses on the multi-view cameras by allowing correspondence across cameras.

While these methods have demonstrated improved performance compared to single-frame approaches, they have their own limitations. Object-centric methods heavily depend on frame-by-frame detection results, and thus, they are susceptible to propagating single-frame errors. Feature-level aggregation methods cannot take full advantage of temporal features due to their naive aggregation methods (*e.g.*, concatenation), and temporal stereo has limited performance on moving objects because of the static scene assumption. Our proposed method overcomes these limitations by introducing prediction into the detection framework and explicitly leveraging motion cues.

3. Method

3.1. Overall Architecture

P2D extends BEVDepth [18] to perform prediction and detection within a single framework. As illustrated in Fig. 2, our proposed P2D consists of a BEV backbone, predic-

tion head, prediction-guided feature aggregation, and detection head. The BEV backbone extracts BEV features from the temporal input images. The BEV prediction head takes BEV features of previous frames as input and predicts object information in the current frame without relying on the current image. The predicted object information and BEV features are then combined into a spatio-temporal feature using the prediction-guided feature aggregation module. Finally, the BEV detection head generates the final detection results by utilizing both the current frame BEV feature and the spatio-temporal feature.

3.2. BEV Backbone

The BEV backbone of P2D consists of an image backbone, a depth network, and a view transformer. The input to the backbone is N multi-view and T multi-timestep images represented as $I = \{I^t \in \mathbb{R}^{N \times H \times W \times 3}, t = 1, 2, \dots, T - 1, T\}$. First, the image backbone (*e.g.*, ResNet [11] with FPN [21]) extracts perspective-view features from the input images. Then, a depth network estimates per-image depth information from these features. Next, a view transformer lifts the perspective-view features into 3D space using the estimated depth and pools them to make BEV representations. The BEV features are represented as $F_{BEV}^{1:T} = \{F_{BEV}^t \in \mathbb{R}^{X_f \times Y_f \times C_f}, t = 1, 2, \dots, T - 1, T\}$, where X_f and Y_f denote the grid size and C_f denotes the channel size of BEV features.

To alleviate the effect caused by ego-motion, we align the coordinates of BEV features from previous images into the current frame, following [12]. For more detailed information on our backbone, please refer to BEVDepth [18].

3.3. BEV Prediction Head

Existing approaches that use temporal images [12, 14, 18, 23] often concatenate all temporal features after alignment to compensate ego-motion. Although simple and intuitive, such a naive strategy can not fully utilize temporal cues, limiting performance gain from temporal frames. Even other approaches like temporal stereo [16, 36, 41] handle temporal features more effectively by enhancing depth estimation with MVS, but they still overlook the importance of motion features. To fully benefit from previous frames, we introduce prediction into the detection framework.

The BEV prediction head uses BEV features only from previous frames to predict object information in the current frame, as follows:

$$P^T = \Phi_p(F_{BEV}^{1:T-1}), \quad (1)$$

where $P^T \in \mathbb{R}^{X_f \times Y_f \times C_o}$ represents predicted object information. C_o is the number of output attributes, including localization, dimension, velocity, orientation of an object, and per-class heatmaps. The per-class heatmaps show the probability of a specific class object in each position of the BEV feature. Φ_p is the detection network, such as the CenterPoint [43] head.

The prediction result P^T provides valuable object-level information, including motion cues, to the downstream network such as the feature aggregation and detection head. It allows the model to leverage explicit object-level motion features. Moreover, supervision on the features from previous frames using the ground truth objects in the current frame can help the model learn the beneficial features of previous frames for detecting the current frames when aggregated. The impact of prediction supervision on the BEV backbone is reported in Table 7.

3.4. Prediction-guided Feature Aggregation

Feature aggregation is a crucial module for effectively merging predicted object information and BEV features. To aggregate BEV features based on predicted object information, we introduce Prediction-guided object queries and Prediction query-based cross attention, as shown in Fig. 3.

Prediction-guided object queries. In contrast to BEV queries in BEVFormer [20], which only has positional information on the BEV space at the initialization stage, we use prediction results P^T as queries to gather temporal features based on predicted object information. However, P^T has a large dimension of $\mathbb{R}^{X_f \times Y_f \times C_o}$, which covers all locations of the BEV space, while objects only occupy a small region of it. Therefore, using P^T as queries in its original form is highly inefficient in terms of computational cost.

To address this problem, we use the object heatmap to select queries. The heatmap represent the probability that objects can exist in a specific space. Specifically, we extract

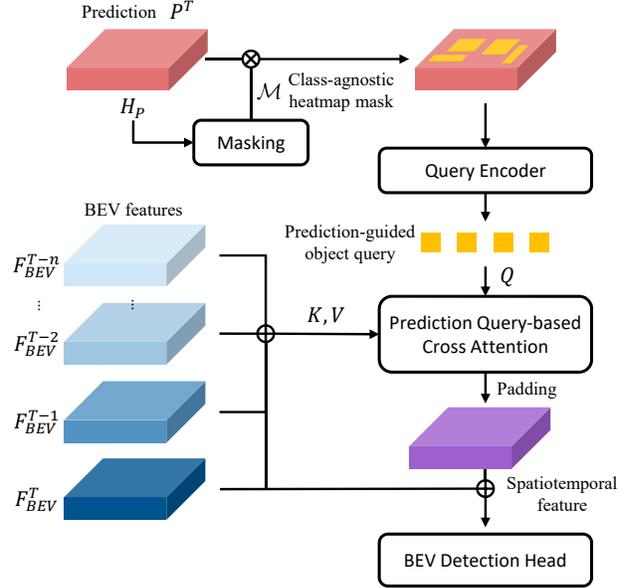


Figure 3. Illustration of the proposed Prediction-guided Feature Aggregation. The class-agnostic heatmap mask is generated from the class heatmap in the prediction results. The query encoder takes the masked prediction results to make the prediction-guided object queries. The prediction query-based cross-attention is a deformable attention module that uses object queries with keys and values from BEV features and fuses the temporal features.

a per-class heatmap represented as $H_P \in \mathbb{R}^{X_f \times Y_f \times N_c}$ from P^T and generate the class-agnostic heatmap represented as $H_{CA} \in \mathbb{R}^{X_f \times Y_f}$ by selecting the maximum probability across all classes as follow:

$$H_{CA}(i, j) = \max_{N_c} H_P(i, j), \quad (2)$$

where (i, j) is the heatmap location index and N_c is the number of object classes.

We then create a class-agnostic query mask $\mathcal{M} \in \mathbb{R}^{X_f \times Y_f}$ by filtering heatmap values with a threshold τ_k .

$$\mathcal{M}(i, j) = \begin{cases} 1 & H_{CA}(i, j) \geq \tau_k, \\ 0 & H_{CA}(i, j) < \tau_k. \end{cases} \quad (3)$$

τ_k in Eq. 3 stands for an adaptive threshold for object probabilities. We choose τ_k as the minimum value of top-k probabilities among H_{CA} so that the number of queries can be fixed. The binary mask indicates the location candidates likely to be occupied by objects.

We apply the query mask to the prediction results to filter out less likely locations. Finally, we embed the masked prediction results into queries using a linear projection.

$$Q = \Phi_q(\mathcal{M} \odot P^T), \quad (4)$$

where $Q \in \mathbb{R}^{K \times C_q}$ stands for the selected prediction-guided object queries, Φ_q is a linear project for query em-

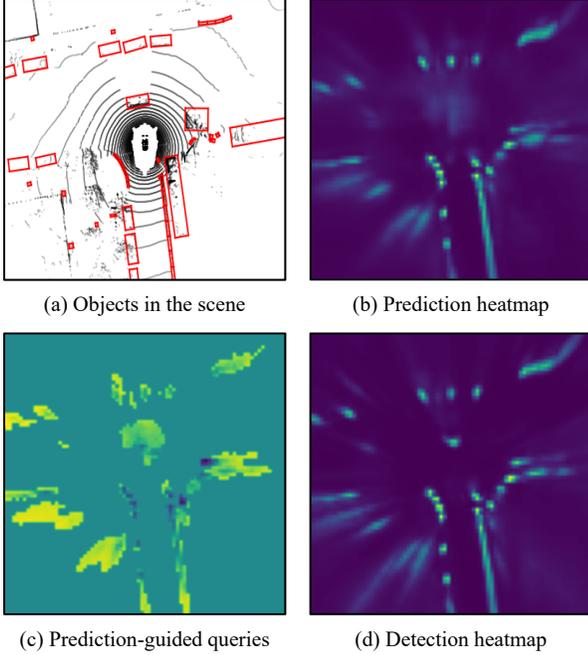


Figure 4. Visualization of the class-agnostic heatmap and prediction-guided object queries. (a) A sample frame with objects in the scenes. (b) A class-agnostic heatmap H_p , which is the output of the prediction head. (c) Prediction-guided object queries generated using the query mask and the predicted object information. (d) The heat map from the final detection results.

bedding, \odot operation denotes element-wise multiplication and K is the number of queries. In this way, we can avoid querying from empty space and reduce computational costs ($K \ll X_f Y_f$). We visualize a sample of the class-agnostic heatmap from both prediction and detection, and masked prediction-guided queries, in Fig. 4.

Prediction query-based cross attention. Temporal features of a moving object are not projected into the same location even after alignment. To effectively aggregate these temporal features, we adopt a deformable attention [46] by setting prediction-guided object queries Q as queries and BEV features $F_{BEV}^{1:T}$ as keys and values. We model the cross attention for temporal BEV features as follows:

$$\begin{aligned}
 & PQCA(Q_p, \{F_{BEV}^{1:T}\}) \\
 &= \sum_{t=1}^T DeformAttn(Q_p, p, F_{BEV}^t + e), \quad (5)
 \end{aligned}$$

where Q_p denote a query located at $p = (i, j)$, respectively. $e = e_s + e^t$ denotes an embedding for positional and temporal dimensions. t is the temporal indexes. Additionally, we utilize zero-padding to match the shape of the output of PQCA with BEV feature's.

We stack each temporal feature level-wise and apply deformable cross attention across all timestamps. Through

this, we can make a spatio-temporal feature by collecting related features in both spatial and temporal dimensions. Our Prediction-guided Feature Aggregation is more effective in modeling the spatio-temporal feature compared to other aggregation methods such as stacking BEV features [12, 18] or Temporal Self-Attention [20]. This is because the prediction-guided object queries provide an object-level prior to the model and work as motion feature-based anchors for the cross-attention mechanism.

The BEV Detection Head concatenates the output of the Prediction Query-based Cross Attention with the BEV feature at the current frame.

$$D^T = \Phi_d(PQCA(Q, \{F_{BEV}^{1:T}\}) \oplus F_{BEV}^T), \quad (6)$$

where D^T is the final detection output and Φ_d is the detection network (e.g., CenterPoint [43] head). It has the same network structure and outputs as BEV Prediction Head but does not share weights.

3.5. Training

P2D is an end-to-end trainable network, and the loss includes two terms: detection loss and prediction loss.

$$\mathcal{L} = \mathcal{L}_{det} + \lambda_p \mathcal{L}_{pred}, \quad (7)$$

where λ_p is balancing weight term. The detection loss \mathcal{L}_{det} consists of classification loss, bounding box loss, and depth estimation loss. Meanwhile, the prediction loss \mathcal{L}_{pred} contains two of them, except for depth estimation loss. We adopt the loss functions as focal loss for classification, L1 loss for bounding box regression, and binary cross-entropy for depth estimation. It is worth noting that P2D does not require any additional annotations.

4. Experiment

4.1. Dataset and Metrics

We conduct experiments on the nuScenes dataset [4], which consists of 1000 videos of around 20 seconds with annotations of 2Hz. The videos are split into three: 700, 150, and 150 scenes for training, validation, and testing, respectively. For the detection task, annotations contain 1.4M 3D bounding boxes of 10 object classes. We adopt the official evaluation metrics to evaluate performance, including nuScenes Detection Score (NDS), mean Average Precision (mAP), mean Average Translation Error (mATE), mean Average Scale Error (mASE), mean Average Orientation Error (mAOE), mean Average Velocity Error (mAVE), and mean Average Attribute Error (mAAE).

4.2. Implementation Details

Unless otherwise specified, we adopt BEVDepth [18] as our baseline model with the ImageNet pretrained ResNet50

Method	Temporal	Backbone	Image Size	mAP \uparrow	NDS \uparrow	mATE \downarrow	mASE \downarrow	mAOE \downarrow	mAVE \downarrow	mAAE \downarrow
PETR [†] [22]		ResNet50	384 × 1056	0.313	0.381	0.768	0.278	0.564	0.923	0.225
BEVDet [†] [13]		ResNet50	256 × 704	0.298	0.379	0.725	0.279	0.589	0.860	0.245
BEVDet4D [12]	✓	ResNet50	256 × 704	0.323	0.453	0.674	0.272	0.503	0.429	0.208
BEVDepth [18]	✓	ResNet50	256 × 704	0.333	0.441	0.683	0.276	0.545	0.526	0.226
BEVStereo [16]	✓	ResNet50	256 × 704	0.344	0.449	0.659	0.276	0.579	0.503	0.216
P2D (BEVDepth)	✓	ResNet50	256 × 704	0.360	0.474	0.643	0.271	0.512	0.412	0.217
P2D (BEVStereo)	✓	ResNet50	256 × 704	0.374	0.486	0.631	0.272	0.508	0.384	0.212
FCOS3D [38]		ResNet101	900 × 1600	0.295	0.372	0.806	0.268	0.511	1.131	0.170
DETR3D [†] [40]		ResNet101	900 × 1600	0.349	0.434	0.716	0.268	0.379	0.842	0.200
PETR [†] [22]		ResNet101	512 × 1408	0.357	0.421	0.710	0.270	0.470	0.885	0.224
UVTR [†] [17]	✓	ResNet101	900 × 1600	0.379	0.483	0.731	0.267	0.350	0.510	0.200
PolarDETR-T [5]	✓	ResNet101	900 × 1600	0.383	0.488	0.707	0.269	0.344	0.518	0.196
BEVDepth* [18]	✓	ResNet101	512 × 1408	0.406	0.490	0.626	0.278	0.513	0.489	0.226
BEVDStereo* [16]	✓	ResNet101	512 × 1408	0.409	0.494	0.651	0.277	0.481	0.451	0.215
BEVFormer [20]	✓	ResNet101	900 × 1600	0.416	0.517	0.673	0.274	0.372	0.394	0.198
P2D (BEVDepth)	✓	ResNet101	512 × 1408	0.420	0.514	0.608	0.268	0.447	0.431	0.212
P2D (BEVStereo)	✓	ResNet101	512 × 1408	0.433	0.528	0.619	0.265	0.432	0.364	0.211
BEVDepth*	✓	ConvNext-B	640 × 1600	0.426	0.521	0.587	0.267	0.393	0.444	0.229
P2D (BEVDepth)	✓	ConvNext-B	640 × 1600	0.460	0.551	0.537	0.259	0.398	0.388	0.212

Table 1. Comparison on the nuScenes *val* set. †: methods with CBGS [45]. *: We reproduce the model without CBGS for a fair comparison.

Method	Backbone	mAP \uparrow	NDS \uparrow	mATE \downarrow	mAOE \downarrow
BEVDepth*	ResNet101	0.396	0.483	0.593	0.533
BEVStereo*	ResNet101	0.404	0.502	0.587	0.518
P2D(BEVDepth)	ResNet101	0.425	0.516	0.549	0.520
P2D(BEVStereo)	ResNet101	0.436	0.530	0.550	0.517

Table 2. Comparison on the nuScenes *test* set. *: We reproduce the model without CBGS for a fair comparison.

[11] backbone and the input image size is resized to 256 × 704. We set the default BEV grid size as 128 × 128. We follow image and BEV data augmentation strategies in [18].

We use two previous frames with 1 second of time interval and set the number of object queries k as 2048. We balance the loss function by setting λ_p as 0.5. We trained the model using AdamW optimizer [26] for 24 epochs with a batch size of 16 on 4 NVIDIA 3090Ti GPUs. The learning rate is set to $2e-4$, and the EMA technique is also used.

4.3. Main Results

We compare our model with existing camera-based detection models on the nuScenes validation dataset [4]. We report the results of P2D with two different baselines: BEVDepth [18] and BEVStereo [16] in Table 1. The evaluation results demonstrate that P2D outperforms other methods and baselines significantly with ResNet50 backbone. Specifically, P2D achieves 2.7% and 3.0% improve-

Baseline	Strategy	mAP \uparrow	NDS \uparrow	mATE \downarrow	mAOE \downarrow	mAVE \downarrow
BEVDepth	P + D	0.360	0.474	0.643	0.512	0.412
	P	0.272	0.422	0.740	0.587	0.294
BEVStereo	P + D	0.374	0.486	0.631	0.508	0.384
	P	0.272	0.418	0.738	0.606	0.308

Table 3. Evaluation of prediction-only results on the nuScenes *val* set. P+D represents the model with both prediction and detection, which is the same as our proposed P2D model. P represents the prediction-only results generated by the BEV Prediction Head, which does not use the current frame.

ment in mAP and 3.3 and 3.7 points improvement in NDS over BEVDepth and BEVStereo, respectively, outperforming other methods with a large margin. In addition, P2D brings a substantial performance boost on velocity estimation, improving by 0.114 m/s and 0.119 m/s (21.5% and 23.7%) in mAVE compared to each baseline. In the case of using a larger backbone and input image size (ResNet101 [11] with 512 × 1408 and ConvNext-B [24] with 640 × 1600), P2D consistently outperforms baselines and other methods both in mAP and NDS.

As shown in Table 2, we also compare the performance with nuScenes test dataset. With the same backbone and image size (ResNet101 with 512 × 1408), P2D still shows improved performance of 2.9% and 3.2% in mAP and 3.3 and 2.8 points in NDS over BEVDepth and BEVStereo, respectively, proving the effectiveness of the proposed method.

Methods	mATE ↓	mASE ↓	mAOE ↓	mAVE ↓	mAAE ↓
BEVDepth	0.815	0.271	0.404	2.010	0.159
+ P2D	0.783	0.256	0.367	1.712	0.149
BEVStereo	0.822	0.269	0.345	1.712	0.149
+ P2D	0.773	0.260	0.243	1.477	0.146

Table 4. Results on the moving objects. Only objects with a velocity higher than 1m/s are evaluated.

PH	PFA	mAP ↑	NDS ↑	mATE ↓	mAOE ↓	mAVE ↓
		0.334	0.448	0.680	0.551	0.469
✓		0.351	0.466	0.668	0.528	0.414
	✓	0.353	0.459	0.662	0.541	0.472
✓	✓	0.360	0.474	0.643	0.512	0.412

Table 5. Ablation study on P2D. PH and PFA denote Prediction Head and Prediction-guided Feature Aggregation, respectively.

Methods	Prev. frames	mAP ↑	NDS ↑	mATE ↓	mAOE ↓	mAVE ↓
BEVDepth	0	0.312	0.357	0.695	0.645	1.144
	1	0.333	0.441	0.683	0.545	0.526
	2	0.334	0.448	0.680	0.551	0.469
	3	0.346	0.451	0.687	0.575	0.461
P2D	2	0.351	0.457	0.672	0.558	0.436
	3	0.362	0.465	0.652	0.434	0.464

Table 6. Experiments on a different number of previous images.

Prediction ability. The quality of the prediction results generated from the prediction head plays a crucial role since it represents the potential of motion features. In addition, it is even more important because P2D uses these prediction results as object queries. Therefore, We evaluated the prediction-only results P^T (the output of the prediction head) and reported them in Table 3 to verify the effectiveness of the prediction. The results show that the prediction-only results can achieve comparable performance to the final detection results, up to 76% and 89% in mAP and NDS, respectively. It proves that the previous frames can estimate objects in the current frame using their motion features, and thus, the motion features can help the model to improve its detection performance.

Moving objects. In the autonomous driving environment, moving objects should be handled more attentively than static objects because moving objects often interact with autonomous agents and can lead to a safety-critical situation. However, previous methods, such as the temporal stereo-based approaches [36, 41] overlook the importance of moving objects and focus on static scenes. To confirm that P2D is advantageous in dynamic scenes, we report the detection results of moving objects in Table 4. In the table, only objects with a ground-truth velocity higher than 1 m/s are evaluated. In both BEVDepth and BEVStereo base-

Prediction supervision	mAP ↑	NDS ↑	mATE ↓	mAOE ↓	mAVE ↓
	0.351	0.457	0.672	0.557	0.464
✓	0.360	0.474	0.643	0.512	0.412

Table 7. Ablation of backbone supervision on prediction loss. The prediction loss does not affect BEV backbone in the model without prediction supervision.

$\lambda(p)$	mAP ↑	NDS ↑	mATE ↓	mASE ↓	mAOE ↓	mAVE ↓	mAAE ↓
0.1	0.352	0.464	0.666	0.279	0.547	0.460	0.212
0.3	0.357	0.463	0.665	0.279	0.545	0.431	0.201
0.5	0.360	0.474	0.643	0.271	0.512	0.412	0.217

Table 8. Ablation of loss balancing weight.

Methods	mAP ↑	NDS ↑	FPS ↑	Memory ↓
BEVDepth	0.334	0.448	8.82	4.26G
P2D	0.360	0.474	10.81	4.40G

Table 9. Comparison of inference time and memory usage.

lines, P2D achieves better performance, especially on the translation (mATE) and velocity error (mAVE). We think that this improvement comes from the prediction scheme in P2D, which explicitly provide motion information by forecasting the location of objects.

4.4. Ablation Studies

We conduct ablation studies to verify the effectiveness of each module and the performance of different hyperparameters. We use the nuScenes *val* set and Table 5 to 9 describe the results of ablation studies.

Prediction head. Table 5 demonstrates the ablation of each module in P2D. The model with only the prediction head concatenates the prediction results and temporal BEV features without an aggregation strategy. Adding Prediction Head to the baseline improves mAP by 1.7 % and NDS by 1.8 points, demonstrating the prediction scheme improves the detection performance. Especially the velocity estimation significantly improves by 11.7%, showing that the prediction strategy helps the model to estimate the motion of objects in the current frame.

Prediction-guided feature aggregation. Adding a deformable attention-based feature aggregation also improves mAP by 1.9% and NDS by 1.1 points, even solely adopted without a prediction head. We hypothesize that our feature aggregation method merges the features of an object along different timesteps, and thus it is beneficial to make features useful. Finally, by combining these two modules, our P2D improves mAP and NDS by 2.6 points, showing the effectiveness of our method.

Number of previous frame. For the fair comparison, we set the number of previous frames as the same and evalu-

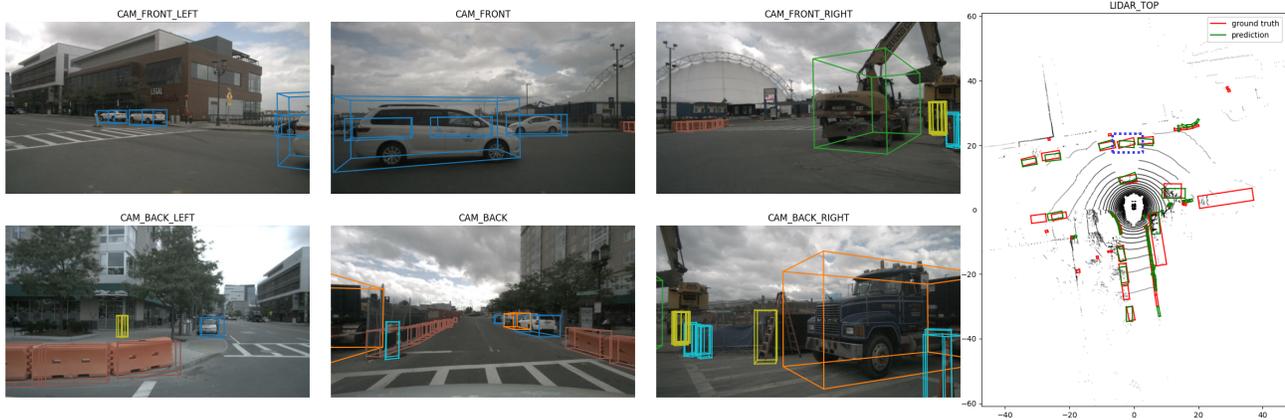


Figure 5. Qualitative results of P2D. The blue dotted rectangle in the BEV view designates the highly occluded object in the image view. Since P2D leverages temporal frames, such an occluded object that appears in the previous frames can be detected.

ated on the nuScenes *val* set. As reported in Table 6, there is still a performance gap between P2D and the baseline BEVDepth with two previous frames. Although there is a significant improvement when a previous frame is used due to the benefit from the multi-frame input, adding another previous frame brings only a marginal improvement, demonstrating that increasing previous frames in a naive manner is barely beneficial. Note that P2D needs at least two previous frames to estimate motion from past frames.

Backbone supervision. P2D has a prediction loss term that provides supervision for the targets in the current frame to the previous frames. We hypothesize that this supervision can guide the backbone to learn how to extract motion-related features from input images. To verify this, we trained P2D with and without the gradient of the prediction loss in the BEV backbone, and Table 7 shows the results. We confirm that with the gradient of the prediction loss on the BEV backbone, the performance improves by 0.9% and 1.7 points in mAP and NDS, respectively, demonstrating the prediction loss makes the BEV backbone learn motion features.

Loss balancing weight. We compare the performance of different values of the loss balancing weight λ_p in Eq. 7. As shown in Table 8, the mAP and velocity estimation improves as the value of λ_p gets larger, proving that the prediction loss helps the model learn motion cues and is beneficial in 3D object detection.

Inference time and memory usage. Table 9 shows the FPS and GPU memory usage during the inference. For a fair comparison, both the baseline and P2D use two previous images with the same backbone and image size. We find that although there is a slight increase in memory usage, P2D runs faster than the baseline by increasing FPS from 8.82 to 10.81.

4.5. Qualitative Results

We visualize a sample case for a qualitative evaluation. P2D is capable of detecting highly occluded objects, as demonstrated by the object enclosed in the blue dotted box in the top view of Fig. 5. Despite being highly occluded in the current frame, this object has been captured in previous frames, enabling P2D to detect it. Additional quantitative results are illustrated in Appendix.

5. Conclusion

In this work, we propose a novel camera-based 3D object detection using temporal images, namely P2D. P2D integrates *Prediction* and *Detection* in a single framework to fully benefit from the sequential images. P2D improves detection performance, including velocity estimation, and we verified that the motion features obtained by prediction are crucial for 3D object detection.

Broader impacts. P2D has shown that utilizing prediction techniques with motion cues can lead to a significant improvement in the performance of 3D object detection models. We believe that further research in this area is warranted to explore how best to leverage prediction strategies for more effective 3D object detection. In addition, P2D has the potential to inspire the development of 3D object tracking models that place a greater emphasis on motion cues and their role in object detection and tracking.

Acknowledgements

This work was supported in part by the Korea Agency for Infrastructure Technology Advancement (KAIA) funded by the Ministry of Land, Infrastructure and Transport and the National Research Foundation of Korea(NRF) funded by Korea Government (Ministry of Science and ICT) under Grants RS-2021-KA162184 and 2022R1A2C200494412.

References

- [1] Gwangbin Bae, Ignas Budvytis, and Roberto Cipolla. Multi-view depth estimation by fusing single-view depth probability with multi-view geometry. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3
- [2] Garrick Brazil and Xiaoming Liu. M3d-rpn: Monocular 3d region proposal network for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 2
- [3] Garrick Brazil, Gerard Pons-Moll, Xiaoming Liu, and Bernt Schiele. Kinematic 3d object detection in monocular video. In *European Conference on Computer Vision*, 2020. 2
- [4] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 5, 6
- [5] Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Chang Huang, and Wenyu Liu. Polar parametrization for vision-based surround-view 3d detection. *arXiv preprint arXiv:2206.10965*, 2022. 6
- [6] Xiaozhi Chen, Kaustav Kundu, Ziyu Zhang, Huimin Ma, Sanja Fidler, and Raquel Urtasun. Monocular 3d object detection for autonomous driving. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 2
- [7] Chia-Chun Cheng and Shang-Hong Lai. 3d object detection from consecutive monocular images. In *Proceedings of the Asian Conference on Computer Vision*, 2020. 2
- [8] Kaiwen Duan, Song Bai, Lingxi Xie, Honggang Qi, Qingming Huang, and Qi Tian. Centernet: Keypoint triplets for object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2019. 2
- [9] Carlos Hernández Esteban and Francis Schmitt. Silhouette and stereo fusion for 3d object modeling. *Computer Vision and Image Understanding*, 2004. 1
- [10] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Ravenstos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020. 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016. 3, 6
- [12] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022. 1, 2, 3, 4, 5, 6
- [13] Junjie Huang, Guan Huang, Zheng Zhu, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. 2, 6
- [14] Yanqin Jiang, Li Zhang, Zhenwei Miao, Xiatian Zhu, Jin Gao, Weiming Hu, and Yu-Gang Jiang. Polarformer: Multi-camera 3d object detection with polar transformer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 1, 2, 4
- [15] Peixuan Li and Jieyu Jin. Time3d: End-to-end joint monocular 3d object detection and tracking for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [16] Yinhao Li, Han Bao, Zheng Ge, Jinrong Yang, Jianjian Sun, and Zeming Li. Bevstereo: Enhancing depth estimation in multi-view 3d object detection with temporal stereo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 1, 3, 4, 6
- [17] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with transformer for 3d object detection. In *Advances in Neural Information Processing Systems*, 2022. 6
- [18] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2023. 1, 2, 3, 4, 5, 6
- [19] Zhuoling Li, Zhan Qu, Yang Zhou, Jianzhuang Liu, Haoqian Wang, and Lihui Jiang. Diversity matters: Fully exploiting depth clues for reliable monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [20] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, 2022. 2, 3, 4, 5, 6
- [21] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017. 3
- [22] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. In *European Conference on Computer Vision*, 2022. 6
- [23] Yingfei Liu, Junjie Yan, Fan Jia, Shuailin Li, Qi Gao, Tiancai Wang, Xiangyu Zhang, and Jian Sun. PetrV2: A unified framework for 3d perception from multi-camera images. *arXiv preprint arXiv:2206.01256*, 2022. 1, 4
- [24] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022. 6
- [25] Zechen Liu, Zizhang Wu, and Roland Tóth. Smoke: Single-stage monocular 3d object detection via keypoint estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020. 2
- [26] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *International Conference on Learning Representations*, 2019. 6
- [27] Yan Lu, Xinzhu Ma, Lei Yang, Tianzhu Zhang, Yating Liu, Qi Chu, Junjie Yan, and Wanli Ouyang. Geometry uncertainty projection network for monocular 3d object detection.

- In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 2
- [28] Arsalan Mousavian, Dragomir Anguelov, John Flynn, and Jana Kosecka. 3d bounding box estimation using deep learning and geometry. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017. 2
- [29] Dennis Park, Rares Ambrus, Vitor Guizilini, Jie Li, and Adrien Gaidon. Is pseudo-lidar needed for monocular 3d object detection? In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 1, 2
- [30] Jinyhung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris Kitani, Masayoshi Tomizuka, and Wei Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. In *International Conference on Learning Representations*, 2023. 1, 3
- [31] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *European Conference on Computer Vision*, 2020. 2
- [32] Zequn Qin and Xi Li. Monoground: Detecting monocular 3d objects from the ground. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2
- [33] Cody Reading, Ali Harakeh, Julia Chae, and Steven L Waslander. Categorical depth distribution network for monocular 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021. 1, 2
- [34] Thomas Roddick, Alex Kendall, and Roberto Cipolla. Orthographic feature transform for monocular 3d object detection. *British Machine Vision Conference*, 2018. 2
- [35] Andrea Simonelli, Samuel Rota Buló, Lorenzo Porzi, Manuel López-Antequera, and Peter Kotschieder. Disentangling monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019. 1
- [36] Tai Wang, Jiangmiao Pang, and Dahua Lin. Monocular 3d object detection with depth from motion. In *European Conference on Computer Vision*, 2022. 1, 3, 4, 7
- [37] Tai Wang, ZHU Xinge, Jiangmiao Pang, and Dahua Lin. Probabilistic and geometric depth: Detecting objects in perspective. In *Conference on Robot Learning*, 2022. 2
- [38] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021. 1, 2, 6
- [39] Yan Wang, Wei-Lun Chao, Divyansh Garg, Bharath Hariharan, Mark Campbell, and Kilian Q Weinberger. Pseudo-lidar from visual depth estimation: Bridging the gap in 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019. 2
- [40] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. In *Conference on Robot Learning*, 2022. 6
- [41] Zengran Wang, Chen Min, Zheng Ge, Yin hao Li, Zeming Li, Hongyu Yang, and Di Huang. Sts: Surround-view temporal stereo for multi-view 3d detection. *arXiv preprint arXiv:2208.10145*, 2022. 1, 3, 4, 7
- [42] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *European Conference on Computer Vision*, 2018. 1, 3
- [43] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021. 4, 5
- [44] 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. *International Conference on Learning Representations*, 2019. 2
- [45] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*, 2019. 6
- [46] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *International Conference on Learning Representations*, 2021. 1, 2, 5