

Progressive Query Refinement Framework for Bird’s-Eye-View Semantic Segmentation from Surrounding Images

Dooseop Choi^{*1,2} and Jungyu Kang¹ and Taeghyun An¹ and Kyoungwan Ahn¹ and KyoungWook Min¹

Abstract—Expressing images with Multi-Resolution (MR) features has been widely adopted in many computer vision tasks. In this paper, we introduce the MR concept into Bird’s-Eye-View (BEV) semantic segmentation for autonomous driving. This introduction enhances our model’s ability to capture both global and local characteristics of driving scenes through our proposed residual learning. Specifically, given a set of MR BEV query maps, the lowest resolution query map is initially updated using a View Transformation (VT) encoder. This updated query map is then upsampled and merged with a higher resolution query map to undergo further updates in a subsequent VT encoder. This process is repeated until the resolution of the updated query map reaches the target. Finally, the lowest resolution map is added to the target resolution to generate the final query map. During training, we enforce both the lowest and final query maps to align with the ground-truth BEV semantic map to help our model effectively capture the global and local characteristics. We also propose a visual feature interaction network that promotes interactions between features across images and across feature levels, thus highly contributing to the performance improvement. We evaluate our model on a large-scale real-world dataset. The experimental results show that our model outperforms the SOTA models in terms of IoU metric. Codes are available at <https://github.com/d1024choi/ProgressiveQueryRefineNet>

I. INTRODUCTION

Perceiving driving environments from surrounding camera images has recently gained significant attention in autonomous driving. This is because camera sensors are not only typically more cost-effective than other sensors, such as LIDAR and RADAR, but also provide rich semantic information that other sensors cannot capture. Detecting 3D objects from the images is one of the actively researched tasks. Another promising task that leverages the abundant information in the images is predicting BEV semantic maps of driving scenes with respect to autonomous vehicles. Since autonomous driving is inherently a geometric problem, where the goal is to navigate a vehicle safely and correctly through 3D space [4], [8], BEV maps can be directly deployed for the subsequent tasks such as motion planning and control.

Presenting input images at multiple resolutions has long been a common practice in various computer vision tasks. This allows NN models to capture both local and global characteristics of the images effectively. One pioneering work is the Feature Pyramid Network (FPN) [15], which is

proposed for the 2D object detection task, and the majority of the models dedicated to the VT leverage the FPN to better capture objects of different sizes in driving images. BEV semantic maps of driving scenes, which typically are created by rendering HD map components and detected 3D objects on a 2D canvas, also possess their own set of global and local characteristics. For instance, 3D objects are depicted as rectangles with varying sizes on the map, whereas lane lines are represented as lines of different lengths. These characteristics can be effectively conveyed when the maps are represented at multiple resolutions.

In this paper, we propose a VT model that leverages MR BEV query maps to better capture the global and local characteristics of driving scenes. In particular, given randomly generated MR BEV query maps, the lowest resolution map is initially updated by a VT encoder. The updated map is then upsampled and merged with a higher resolution one to be fed into a subsequent VT encoder for further updates. This process repeats until the updated query map reaches the target resolution. The final query map is then obtained by adding the lowest resolution to the target resolution. We supervise both the lowest and final maps during training. As a consequence, the VT encoder dedicated to the lowest resolution can learn to capture the global characteristics while the other VT encoders dedicated to higher resolution ones contribute to capturing missing details. We also propose a visual feature interaction network to promote the interaction between features across images and feature levels, which in turn contributes to the performance improvement.

II. RELATED WORKS

View Transformation: Transforming features in Camera View (CV) to BEV is considered an ill-posed problem, primarily because camera depth information is generally assumed to be unknown. To address this issue, some early works proposed training NNs to directly transform image features to features in BEV space without considering the camera pose and depth [20], [24]. LSS [23] is the first to propose predicting the depth information for the forward projection, where image features are projected into BEV space according to the predicted depth and camera intrinsic/extrinsic parameters. Based upon the forward projection paradigm, subsequent works proposed methods that better capture the BEV representation for tasks such as BEV segmentation [8], [9], [19], [33], motion prediction and planning [1], [8], [9], and 3D object detection [19], [30].

The backward projection paradigm projects 3D reference points in BEV space to image plane to sample image features

*Corresponding author

¹D. Choi, T. An, K. Ahn, and K. Min are with Superintelligence Creative Research Laboratory, ETRI, South Korea {d1024.choi, tekkeni, mobileguru, kwmin92}@etri.re.kr

²D. Choi is Faculty of Artificial Intelligence, University of Science and Technology, South Korea

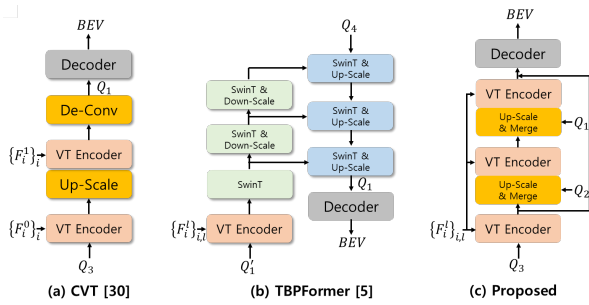


Fig. 1: Different progressive query refinement methods. (c) briefly depicts the proposed architecture.

and back-project them to the BEV grid. Consequently, it avoids the additional operations for the depth prediction and mitigates the sparse mapping problem in the forward projection. OFT [25] is the first to propose using the backward projection for the VT. Building upon this, BEVFormer [13] further suggested leveraging the powerful capability of the deformable attention [32] for learning the VT. In accordance with the backward projection paradigm, subsequent works proposing improved performance have been introduced for various tasks such as BEV segmentation [6], 3D segmentation [10], and 3D object detection [27], [29]. In this paper, we deploy the VT encoder proposed in BEVFormer [13] for the BEV query map update. However, as illustrated in Fig. 1, any query-based VT encoder can be easily deployed within our framework.

Multi-Resolution BEV Query Representation: Representing the BEV space with MR query maps has been adopted in preceding works [17], [21], [30], [31], primarily for the reduction of computational complexity. Specifically, as depicted in Figure 1a, an LR BEV query map is updated using a VT encoder and then upsampled for further updates in a subsequent VT module. This process is repeated until the resolution reaches an affordable level. The BEV feature map of the target resolution is then finally obtained via an up-sampling network. TBPFormer [6] proposed first updating the target resolution query map via BEVFormer [13] and then producing MR query maps from the target resolution map through SwinT [18] and downsampling operations. The generated MR query maps are then utilized to progressively update and upsample another LR query map up to the final resolution as depicted in Figure 1b.

In this paper, we begin with a set of randomly generated MR query maps. These are progressively updated and merged starting from the lowest resolution to the target resolution as depicted in Figure 1c. The final query map is then obtained by adding the lowest resolution to the target resolution. While the existing works supervise only the target resolution during training, we propose supervising both the lowest and final maps. Supervising the lowest resolution map during the training forces the VT encoders dedicated to the higher resolution maps to learn missing details in the lowest map, which can be regarded as the residual learning [7].

Visual Feature Interaction : Objects and backgrounds

in an image tend to have specific relationships within a context. For example, pedestrians tend to be on sidewalks, while vehicles are on roads. In order to capture these relationships and leverage them for the tasks at hand, many works have proposed promoting interactions between features in an image through attention mechanism [3], [5], [18], [28], [32]. For the VT task, however, less attention has been devoted to promoting the interactions primarily because BEV queries can easily attend to relevant regions across different images through the attention, and interactions arise implicitly during the attention process. Peng et al. [22] proposed promoting interactions between features in an image at different levels through deformable attention [32]. Pan et al. [21] attempted to promote the interactions implicitly through their proposed bi-direction cross-attention in which image features are updated from the refined BEV query maps. In this paper, we propose a visual feature interaction network designed to promote interactions between features not only across feature levels but also across images.

III. PROPOSED APPROACH

A. Problem Formulation

Suppose that an autonomous vehicle (AV) is equipped with N_c surrounding cameras. The i -th camera image acquired at time t is denoted as $I_{i,t} \in \mathbb{R}^{H_I \times W_I \times 3}$, where H_I and W_I denote the height and width of the image, respectively. Our target is to obtain a BEV feature map $\mathbf{B}_t \in \mathbb{R}^{X \times Y \times C}$, which is utilized as input to subsequent NNs for the semantic BEV map prediction, from the images $\{I_{i,t}\}_{i=1}^{N_c}$. Here, X and Y respectively denote the length and width of the BEV grid space spanned by the vehicle pose at t , and C denotes the BEV feature dimension. In the rest of this paper, we omit t for the sake of simplicity. Our proposed VT module takes a set of MR BEV query maps $\{\mathbf{Q}_s\}_{s=1}^{N_s}$ and MR image feature maps $\{\mathbf{F}_i^l\}_{i=1, l=1}^{N_c, N_l}$ as input and produces \mathbf{B} by updating and merging the query maps progressively. Here, N_s is the number of the MR query maps and $\mathbf{Q}_s \in \mathbb{R}^{\frac{X}{2^{s-1}} \times \frac{Y}{2^{s-1}} \times C}$. The image feature map \mathbf{F}_i^l is the l -th layer output of an image backbone (e.g., ResNet [7]) when I_i is used as input.

B. Network Architecture

Figure 2 briefly illustrates the overall architecture of our proposed BEV segmentation network. First, MR feature maps $\{\mathbf{F}_i^l\}$ are extracted from $\{I_i\}$ using a backbone network. The proposed visual feature interaction (VFI) module is then applied to the MR feature maps to promote interactions between the features both across levels and across images. Next, the proposed VT module, which consists of VT encoders and upsamplers, updates and merges $\{\mathbf{Q}_s\}$ progressively to produce $\mathbf{B} = \mathbf{Q}_1 + \text{UpScale}(\mathbf{Q}_{N_s}) \in \mathbb{R}^{X \times Y \times C}$. Finally, a class header network, which is dedicated to a specific object or road element class, is applied to \mathbf{B} and the final BEV map $\mathbf{BEV} \in \mathbb{R}^{X \times Y}$ is predicted using a subsequent map decoder network. As we mentioned, supervising both \mathbf{B} and \mathbf{Q}_{N_s} during training can be regarded as residual learning

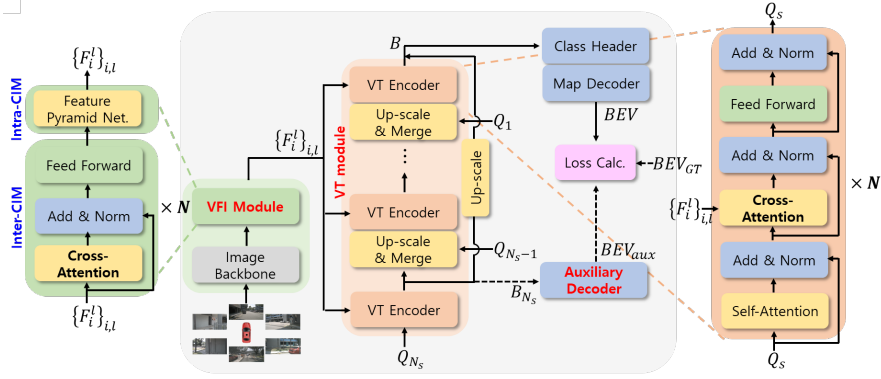


Fig. 2: The overall architecture of the proposed segmentation network




Fig. 3: Examples of (a) the conventional and (b) the proposed 2D reference points patterns.

[7], resulting in a significant improvement in segmentation performance.

Visual feature interaction (VFI) module is devised to promote interactions between the MR image features. It comprises two sub-modules: 1) Intra-camera interaction module (Intra-CIM), 2) Inter-camera interaction module (Inter-CIM). Intra-CIM is introduced to promote interactions between features located at the same pixel position across different feature levels and is implemented by the FPN [15]. Inter-CIM is introduced to promote interactions between features across different camera images. Given a set of image features *at the same feature level*, $\{\mathbf{F}_i\}$, the interactions are promoted through our proposed multi-head deformable attention as follows:

$$f_{DA}(\mathbf{F}_i(\mathbf{p}), \{\mathbf{F}_j\}_{j=1}^{N_c}) = \sum_{h \in \mathcal{H}} \sum_{\mathbf{p} \in \mathcal{P}_h} \sum_{i=1}^{N_c} \alpha_{\mathbf{p},i}^h \mathbf{F}_i(\mathbf{p} + \Delta \mathbf{p}_i^h), \quad (1)$$

where \mathcal{H} is a set of head indices, \mathcal{P}_h is a set of 2D reference points for a head of index h , and $\alpha_{\mathbf{p},i}^h$ denotes an attention weight and satisfies $\sum_{\mathbf{p},i} \alpha_{\mathbf{p},i}^h = 1$. Our deformable attention distinguishes itself from the original [32] through the design of the reference points pattern and the calculation of offsets $\{\Delta \mathbf{p}_i^h\}$ and weights $\{\alpha_{\mathbf{p},i}^h\}$. Let us elaborate on our reference points pattern using Fig. 3b. For the deformable attention operation, we assign each head four reference points, covering specific areas in an image. We illustrate in Fig. 3b the reference points for each head with different colors. As a consequence, with the eight heads having distinct reference points, which cover different image areas, our model can easily locate the image regions to be attended to. On the

other hand, in the conventional design, the reference points for each head are the same as the pixel position of the query feature, \mathbf{p} , as seen in Fig. 3a.

In the conventional deformable attention [32], the offsets and weights are predicted directly from $\mathbf{q}_{\text{emb}} = \mathbf{F}_i(\mathbf{p}) + \mathbf{P}(\mathbf{p})$ through MLPs, where $\mathbf{P}(\mathbf{p})$ denotes a sinusoidal positional embedding for the pixel position \mathbf{p} . We, in this paper, predict the offsets and weights as proposed in [32] with the following two modifications: 1) We introduce trainable embedding vectors $\{\mathbf{c}_i\}_{i=1}^{N_c}$ to encourage the VFI module to distinguish image features across different images and to enhance the prediction of the offsets and weights. As a result, the offsets and weights are predicted from $\mathbf{q}_{\text{emb}} = \mathbf{F}_i(\mathbf{p}) + \mathbf{P}(\mathbf{p}) + \mathbf{c}_i$ through MLPs. 2) We limit the maximum allowable magnitude of the offsets to ensure that a point $\mathbf{p} + \Delta \mathbf{p}_i^h$ can cover a specific area of an image as follows:

$$\Delta \mathbf{p}_i^h = \delta \cdot \text{Tanh}(\text{MLP}(\mathbf{q}_{\text{emb}})), \quad (2)$$

where $\text{Tanh}()$ denotes the hyperbolic tangent function. δ is a positive constant and is set to 0.25 in this paper.

View Transformation (VT) Encoder repeatedly updates \mathbf{Q}_s using $\{\mathbf{F}_i^l\}$. We deploy the Transformer encoder proposed in [13], which benefits from the deformable attention [32] in terms of complexity and training speed. Specifically, let \mathbf{V} be the projection of \mathbf{Q}_s through a MLP. Then, the self-attention module in the VT encoder updates \mathbf{Q}_s via the following deformable attention:

$$f_{DA}(\mathbf{Q}_s(\mathbf{p}), \mathbf{V}) = \sum_{h \in \mathcal{H}} \sum_{z=1}^Z \alpha_z^h \mathbf{V}(\mathbf{p} + \Delta \mathbf{p}_z^h), \quad (3)$$

where $\mathbf{Q}_s(\mathbf{p})$ denotes the element of \mathbf{Q}_s at the spatial location \mathbf{p} in the BEV grid space. Z is the number of the value sampling operation. By adding the result of (3) to $\mathbf{Q}_s(\mathbf{p})$ and applying a feed forward network (FFN) to the addition, the interaction between queries in \mathbf{Q}_s is promoted.

The cross-attention module in the VT encoder updates \mathbf{Q}_s using the visual features $\{\mathbf{F}_i^l\}$ through the deformable

Model	Input Size	Vehicle	Pedestrian	Drivable Area	Lane Lines
CVT [31]	224 × 480	36.0	-	74.3	-
LSS [23]	224 × 480	32.1	15.0	72.9	20.0
FIERY- <i>s</i> [8]	224 × 480	35.8	-	-	-
ST-P3 [9]	224 × 480	40.1	14.5	75.9	-
TBPFormer- <i>s</i> * [6]	224 × 480	43.6	16.1	-	-
BAEFormer [6]	224 × 480	38.9	-	76.0	-
BEVFormer- <i>s</i> [13]	900 × 1600	43.2	-	80.7	21.3
Ours	224 × 480	43.7	15.7	81.1	25.6

TABLE I: Experiments on nuScenes dataset. The values in this table are from the corresponding papers. * denotes that values are estimated from the respective paper assuming no data augmentation during training. *s* stands for *static*, indicating that only current-time image data is used for prediction. We exclude the result of ST-P3 [9] for *Lane Lines* since it didn't follow the common practice for the ground-truth BEV semantic map generation.

attention defined as follows:

$$\begin{aligned}
& f_{DA}^h(\mathbf{Q}_s(\mathbf{p}), \{\mathbf{F}_i^l\}, \{\mathcal{T}_i\}) \\
&= \frac{1}{|\mathcal{V}_{hit}|} \sum_{i \in \mathcal{V}_{hit}} \sum_{z=1}^Z \sum_{l=1}^{N_l} \alpha_{z,l}^i \mathbf{V}_i^l(\mathcal{T}_i(\mathbf{p}, i, z) + \Delta \mathbf{p}_{z,l}^i), \quad (4)
\end{aligned}$$

where f_{DA}^h denotes the deformable attention dedicated to the h -th header and \mathbf{V}_i^l is the projection of \mathbf{F}_i^l through a MLP. \mathcal{T}_i is the transformation function that projects a 3D point in the BEV space into a 2D point in the i -th image space. Consequently, $\mathcal{T}_i(\mathbf{p}, i, z)$ is a 2D point in the i -th image space projected from the z -th 3D reference point defined on the position \mathbf{p} in the BEV space. \mathcal{V}_{hit} is the set of image indices where at least one of the 3D reference points is projected to. Finally, we note that the attention weight α and offset $\Delta \mathbf{p}$ are predicted directly from $\mathbf{Q}_s(\mathbf{p})$ through MLPs after a positional sinusoidal embedding is added to $\mathbf{Q}_s(\mathbf{p})$.

Query Map Merge Module fuses a previously updated LR query map \mathbf{Q}_{s-1} with a higher resolution query map \mathbf{Q}_s for a subsequent query update. We explored diverse options for the fusion and empirically found that a simple addition operation yields the best performance. Specifically, \mathbf{Q}_{s-1} first undergoes bilinear interpolation to match the size of \mathbf{Q}_s and then is added to \mathbf{Q}_s . Finally, \mathbf{Q}_s goes through a subsequent VT encoder.

Auxiliary Task is introduced to encourage our model to capture the global characteristics of driving scene more effectively. Our auxiliary decoder depicted in Figure 2 progressively increases the resolution of the lowest resolution query map up to the target resolution to obtain the final output $\mathbf{BEV}_{aux} \in \mathbb{R}^{X \times Y}$. It comprises a series of NN modules, each of which consists of Up-scaler, Conv, BN [12], and Relu. During the training, \mathbf{BEV}_{aux} is forced to match the ground-truth BEV semantic map, \mathbf{BEV}_{GT} . As shown in the section IV, our model achieves the best performance when only the lowest resolution query map undergoes the auxiliary task.

C. Losses

To train our model, we minimize the final loss \mathcal{L} defined as follows:

$$\mathcal{L} = \sum_{c \in \mathcal{C}} \lambda_c (\mathcal{L}_{main}^c + \alpha \mathcal{L}_{aux}^c), \quad (5)$$

where \mathcal{L}_{main}^c and \mathcal{L}_{aux}^c denote the focal losses [16] for a specific class c calculated from \mathbf{BEV} , \mathbf{BEV}_{aux} , and \mathbf{BEV}_{GT} . α and λ_c are pre-defined constants. α balances the contributions of \mathcal{L}_{main} and \mathcal{L}_{aux} to the final loss. λ_c balances the contributions of each class to the final loss.

IV. EXPERIMENTS

A. Dataset and Evaluation Settings

A large-scale real-world dataset, nuScenes [2], is used to evaluate our model. nuScenes is a collection of 1000 scenes acquired over diverse weather, time-of-day, and traffic conditions. Following the previous works [6], [13], [31], 28,130 and 6,019 key frames are used for the training and test sets, respectively. To generate the ground-truth BEV semantic maps at a resolution of (200 × 200), 3D bounding boxes of vehicles and pedestrians or HD map elements in 100m × 100m area around the ego-vehicle are orthographically projected onto the ground plane, following the standard practice [8], [23]. For an evaluation metric, we use Intersection-over-Union (IoU) score between the ground-truth and its prediction.

B. Network and Training Details

We use as an image backbone ResNet-50 [7] pre-trained on ImageNet [26]. Three consecutive mid-layer feature maps ($N_l = 3$), each of which corresponds to down-scaling factors of ×4, ×8, and ×16 from the original size, are used as the visual feature maps. The feature maps further undergo a shallow CNN to match the channel dimension with that of the query maps. For the class headers, we use shallow CNNs consisting of Conv, BN [12], and Relu. For the map decoders, we deploy the mask decoder [14]. Our model is trained with a batch size of 4 for 28 epochs. We optimize our model using AdamW [11] optimizer with learning rate $2e^{-4}$ and weight decay $1e^{-7}$. We also set $N_s = 3$ and $\alpha = 1$ in Eqn. 5. Lastly, in accordance with common practice [21],



Fig. 4: Visualization of the prediction results. The first, second, and third columns respectively are the surrounding images, the ground-truth BEV semantic map, and the prediction. *Vehicle*, *Pedestrian*, *Drivable Space*, and *Lane line* are color-coded with orange, blue, grey, and red, respectively.




Fig. 5: Visualization of the ground-truth BEV map and its predictions. The first, second, and third columns respectively correspond to the ground-truth, the prediction from the lowest resolution query map, and the final prediction. The green arrows highlight how missing details are restored through our progressive refinement framework.

[31], we trained a separate NN for each class to achieve the results presented in the subsequent sections. This separation is adopted due to the potential underperformance of jointly trained models, a phenomenon known as *negative transfer* in multi-task learning [13].

C. Subjective and Objective Results

We objectively compare our model with the SOTA models in Table I. Since our model does not consider previously obtained image data for the current prediction, we report the prediction performance of the *static* version of the SOTA models in the table for fair comparisons. We can see in the table that our proposed model achieves the best performance across almost all class categories. Specifically, our model outperforms BEVFormer-s [13] despite the input image size for our model being four times smaller than that for BEVFormer-s [13]. This result demonstrates that our model

Model	Vehicle	Drivable Space
M1	42.1	79.6
M2	42.3	79.8
M3	42.8	80.3
M4	42.0	79.0
M5	43.3	80.3
M6	43.2	81.0
Ours	43.7	81.1

TABLE II: Ablation study conducted on nuScenes

effectively captures objects of various sizes in driving scenes. TBPFormer-s [6] also shows the performance comparable to ours. However, it has 1.6 times more trainable parameters than ours (117.9M v.s. 73.7M), which mainly originates from its repeated deployments of SwinT [18] for the hierarchical query refinement.

In Figure 4, we visualize the prediction results of our model alongside the ground-truth. We can see in the figure that our model captures the shape of the dynamic road agents and static road elements well. Specifically, our model effectively distinguishes vehicles gathered owing to our query refinement framework and its dedicated training method. Figure 5 further demonstrates the effectiveness of the framework.

D. Ablation Study

In Table II, we demonstrate the effectiveness of the VFI module and our progressive refinement framework in terms of the prediction performance. In the table, **M1** and **M2** denote the proposed model without the VFI module and with the Inter-CIM removed, respectively. For **M3**, we replace the proposed deformable attention for the Inter-CIM with the conventional one. The table shows that the proposed deformable attention enhances the interaction between the features across images, leading to a significant improvement

in prediction performance compared to the conventional attention (**M3** v.s. **Ours**). **M4** and **M5** respectively denote the proposed model with a single query map ($N_s = 1$) and MR query maps ($N_s = 3$) without applying the auxiliary task during the training. In contrast, for **M6** ($N_s = 3$), we exclude the addition of the lowest resolution map to the highest one and apply the auxiliary task for all query maps except for the final resolution. The table and Figure 5 demonstrate that the proposed progressive query refinement framework, along with its dedicated training (the auxiliary task), enhances our network’s ability to capture the global and local characteristic of driving scenes.

V. CONCLUSIONS

This paper proposes a new NN architecture that progressively updates and merges query maps from the lowest resolution to the final resolution for BEV semantic segmentation. Our model’s ability to capture the global and local characteristics of driving scenes is significantly improved through the proposed residual learning. In addition, a new visual feature interaction method is proposed to further enhance the prediction performance of our model. Our future research plan includes incorporating temporal information such as previously obtained image data into our framework.

ACKNOWLEDGMENT

This work was supported by IITP grant funded by the Korea government (MSIT) (RS-2023-00236245, Development of Perception/Planning AI SW for Seamless Autonomous Driving in Adverse Weather/Unstructured Environment)

REFERENCES

- [1] A. K. Akan and F. Guney, “Stretchbev: stretching future instance prediction spatially and temporally,” in *Eur. Conf. Comput. Vis.*, 2022.
- [2] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, “nuscenes: a multimodal dataset for autonomous driving,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020.
- [3] N. Carion, F. Massa, G. Synnaeve, N. Usunier, A. Kirillov, and S. Zoraruyko, “End-to-end objection detection with transformers,” in *Eur. Conf. Comput. Vis.*, 2020.
- [4] D. Choi, S.-J. Han, K.-W. Min, and J. Choi, “Pathgan: Local path planning with attentive generative adversarial networks,” *ETRI Journal*, vol. 44, pp. 1004–1019, 2022.
- [5] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: transformers for image recognition at scale,” in *Int. Conf. on Learn. Represent.*, 2020.
- [6] S. Fang, Z. Wang, Y. Zhong, J. Ge, S. Chen, and Y. Wang, “Tbp-former: learning temporal bird’s-eye-view pyramid for joint perception and prediction in vision-centric autonomous driving,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023.
- [7] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016.
- [8] A. Hu, Z. Murez, N. Mohan, S. Dudas, J. Hawke, V. Badrinarayanan, R. Cipolla, and A. Kendall, “Fiery: future instance prediction in bird’s-eye view from surround monocular cameras,” in *Int. Conf. Comput. Vis.*, 2021.
- [9] S. Hu, L. Chen, P. Wu, H. Li, J. Yan, and D. Tao, “St-p3: end-to-end vision-based autonomous driving via spatio-temporal feature learning,” in *Eur. Conf. Comput. Vis.*, 2022.
- [10] Y. Huang, W. Zheng, Y. Zhang, J. Zhou, and J. Lu, “Tri-perspective view for vision-based 3d semantic occupancy prediction,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023.
- [11] F. H. I. Loshchilov, “Decoupled weight decay regularization,” in *Int. Conf. on Learn. Represent.*, 2018.
- [12] S. Ioffe and C. Szegedy, “Batch normalization: accelerating deep network training by reducing internal covariate shift,” in *Int. Conf. on Machine Learn.*, 2015.
- [13] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai, “Bevformer: learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers,” in *Eur. Conf. Comput. Vis.*, 2022.
- [14] Z. Li, W. Wang, E. Xie, Z. Yu, A. Anandkumar, J. M. Alvarez, P. Luo, and T. Lu, “Panoptic segformer: delving deeper into panoptic segmentation with transformers,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022.
- [15] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie, “Feature pyramid networks for object detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2017.
- [16] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar, “Focal loss for dense object detection,” in *Int. Conf. Comput. Vis.*, 2017.
- [17] Y. Liu, T. Wang, X. Zhang, and J. Sun, “Petr: position embedding transformation for multi-view 3d object detection,” in *Eur. Conf. Comput. Vis.*, 2022.
- [18] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, “Swin transformer: hierarchical vision transformer using shifted windows,” in *Int. Conf. Comput. Vis.*, 2021.
- [19] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. L. Rus, and S. Han, “Bevfusion: multi-task multi-sensor fusion with unified bird’s-eye view representation,” in *IEEE Int. Conf. Robotics and Automation*, 2023.
- [20] B. Pan, J. Sun, H. Y. T. Leung, A. Andonian, and B. Zhou, “Cross-view semantic segmentation for sensing surroundings,” *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4867–4873, 2020.
- [21] C. Pan, Y. He, J. Peng, Q. Zhang, W. Sui, and Z. Zhang, “Baeformer: bi-directional and early interaction transformers for bird’s eye view semantic segmentation,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023.
- [22] L. Peng, Z. Chen, Z. Fu, P. Liang, and E. Cheng, “Bevsegformer: Bird’s eye view semantic segmentation from arbitrary camera rigs,” in *2023 IEEE/CVF Winter Conf. on Appl. of Comput. Vis. (WACV)*, 2023.
- [23] J. Philion and S. Fidler, “Lift, splat, shoot: encoding images from arbitrary camera rigs by implicitly unprojecting to 3d,” in *Eur. Conf. Comput. Vis.*, 2020.
- [24] T. Roddick and R. Cipolla, “Predicting semantic map representations from images using pyramid occupancy networks,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2020.
- [25] T. Roddick, A. Kendall, and R. Cipolla, “Orthographic feature transform for monocular 3d object detection,” in *the British Machine Vis. Conf.*, 2019.
- [26] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein, “Imagenet large scale visual recognition challenge,” in *arXiv:1409.0575*, 2014.
- [27] Y. Wang, Y. Chen, and Z. Zhang, “Frustumformer: adaptive instance-aware resampling for multi-view 3d detection,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023.
- [28] Z. Xia, X. Pan, S. Song, L. E. Li, and G. Huang, “Vision transformer with deformable attention,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022.
- [29] C. Yang, Y. Chen, H. Tian, C. Tao, X. Zhu, Z. Zhang, G. Huang, H. Li, Y. Qiao, L. Lu, J. Zhou, and J. Dai, “Bevformer v2: adapting modern image backbones to bird’s-eye-view recognition via perceptual supervision,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2023.
- [30] W. W. A. A. T. L. J. M. A. Z. Li, Z. Yu, “Fb-bev: bev representation from forward-backward view transformations,” in *Int. Conf. Comput. Vis.*, 2023.
- [31] B. Zhou and P. Krahenbuhl, “Cross-view transformers for real-time map-view semantic segmentation,” in *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022.
- [32] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, “Deformable detr: deformable transformers for end-to-end object detection,” in *Int. Conf. on Learn. Represent.*, 2021.
- [33] X. Zhu, V. Zyrianov, Z. Liu, and S. Wang, “Mapprior: bird’s-eye view map layout estimation with generative models,” in *Int. Conf. Comput. Vis.*, 2023.