

Uncertainty-Instructed Structure Injection for Generalizable HD Map Construction

Xiaolu Liu^{1*}, Ruizi Yang^{1*}, Song Wang¹, Wentong Li², Junbo Chen^{3†}, Jianke Zhu^{1†}
¹Zhejiang University ²NCAA ³Udeer.ai

{xiaoluliu, ruiziyang, songw, jkzhu}@zju.edu.cn, wentong_li@ncaa.edu.cn, junbo@udeer.ai

Abstract

Reliable high-definition (HD) map construction is crucial for the driving safety of autonomous vehicles. Although recent studies demonstrate improved performance, their generalization capability across unfamiliar driving scenes remains unexplored. To tackle this issue, we propose **UIGenMap**, an uncertainty-instructed structure injection approach for generalizable HD map vectorization, which concerns the uncertainty resampling in statistical distribution and employs explicit instance features to reduce excessive reliance on training data. Specifically, we introduce the perspective-view (PV) detection branch to obtain explicit structural features, in which the uncertainty-aware decoder is designed to dynamically sample probability distributions considering the difference in scenes. With probabilistic embedding and selection, UI2DPrompt is proposed to construct PV-learnable prompts. These PV prompts are integrated into the map decoder by designed hybrid injection to compensate for neglected instance structures. To ensure real-time inference, a lightweight Mimic Query Distillation is designed to learn from PV prompts, which can serve as an efficient alternative to the flow of PV branches. Extensive experiments on challenging geographically disjoint (geo-based) data splits demonstrate that our UIGenMap achieves superior performance, with +5.7 mAP improvement on the nuScenes dataset. Source code is available at <https://github.com/xiaoluliu2/UIGenMap>.

1. Introduction

Reliable HD map construction is crucial for driving safety of autonomous vehicles, which provides essential geometric information for ego-localization, perceptions, path planning, and predictions [10, 13, 23, 45]. Recently, online high-definition (HD) map vectorization has attracted considerable attention for its cost-effectiveness and ability to

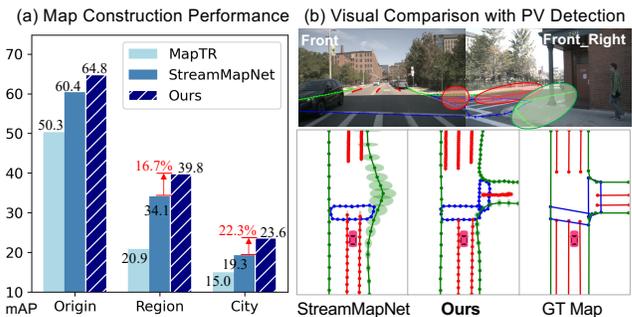


Figure 1. (a) Comparisons on original and different geo-based data partitions among previous methods and ours. (b) Visual comparisons with circles on BEV map to represent the learned uncertainties, in which PV instances serve as compensation and provide a detailed understanding of map perception with generalization.

provide real-time updates [25, 31, 56]. Given images captured by onboard cameras, current map construction relies on transformer scheme [57] to convert perspective-view (PV) image features into bird’s-eye-view (BEV) space. Then, a decoder follows to predict the map elements.

Despite the achievements on widely adopted benchmarks [4, 26, 28], little attention has been paid to investigating the generalization capabilities, which refer to the model’s performance on unfamiliar driving scenarios that are different from the training sets. As discussed in [27, 50], in the widely used public datasets [1, 46], there are a large number of geographical overlaps between the training and validation scenes. The models are then more likely to memorize similar driving scenes rather than fully capture their underlying road structures. Variations in driving scenes, such as the contrast between bustling city centers and suburbs with simple road structures, highlight significant shifts in feature distributions and reveal the model’s limitations in robustness. Besides, learning-based BEV transformation will inevitably lead to geometric errors and loss of texture details. As shown in Figure 1(a), previous methods [25, 50] exhibit notable performance degradation on region-based and city-based data splits under changing driving scenes.

To tackle the above issue, we propose the innovative em-

*Equal Contribution

†Corresponding authors

ployment of model uncertainty targeted for more generalizable HD map constructions. On the one hand, the uncertainty estimation process involves learning both the statistical mean and variance value, enabling dynamic resampling based on probabilistic distributions. Given the discrepancy in feature distribution across different driving scenes, such a statistical resampling strategy enhances the model’s dynamic adaptability in unfamiliar environments. However, the estimated model uncertainty quantifies the confidence in its predictions, which can serve as a reliable indicator for further feature selection and enhancement. As illustrated in Figure 1(b), direct 2D perception can capture more intuitive semantic and angular structural information, which might be missed or incorrectly learned during the implicit PV-to-BEV conversions. Thus, combined with uncertainty-based feature selection, more reliable explicit information can be emphasized as compensation to boost the robustness of BEV-level map construction.

Based on the above discussion, in this work, we propose *UIGenMap*, an uncertainty-instructed PV structure injection strategy for more generalizable HD map vectorization. Specifically, a PV detection branch is introduced to provide explicit structural instances. In both PV and BEV branches, we move beyond the vanilla transformer structure by designing the statistically sensitive uncertainty-aware decoder (UA-Decoder), which constructs the probabilistic outputs by uncertainty-aware attention and finer point-level uncertainty head. Leveraging the PV uncertainty output, we introduce the UI2DPrompt to construct learnable PV prompts. These PV prompts are integrated into the BEV decoder by the designed hybrid injection mechanism, which acts as structural compensation for the BEV map prediction. Consequently, reliable structural representation from PV space can be injected to enhance the model’s capability, thereby improving its perceptual ability in complex environments. Considering the demands for real-time inference, we also design a lightweight mimic query distillation module to learn from constructed PV prompts. Then during inference, the mimicked query can replace the PV detection flow to reduce extra computational burdens.

To prioritize the evaluation of generalization capability, we assess our method using geographically disjoint (geo-based) data partitions on the nuScenes [1] and Argoverse2 [46] datasets. Extensive experiments demonstrate that our *UIGenMap* achieves state-of-the-art performance with higher adaptive ability, especially with +5.7 mAP improvement on the region-based partition and +4.3 mAP on the more challenging city-based split on nuScenes. The main contributions can be summarized as follows:

- We propose *UIGenMap*, an uncertainty-instructed approach to more generalizable HD map vectorization with explicit structure injection.
- The UA-Decoder and UI2DPrompt are proposed to con-

struct structural PV prompts, which are dynamically integrated with BEV perception by hybrid injection.

- Lightweight MQ-Distillation design supports real-time inference without the additional computational overhead of the PV detection branches.
- Promising results on more challenging geo-based data partitions show that *UIGenMap* outperforms previous approaches and achieves powerful generalization capability.

2. Related Work

Vectorized HD Map Construction. Conventional HD map construction approaches [38, 39, 58] rely on SLAM-based annotations, which require tedious offline manual labeling at high costs. Recent studies [6, 11, 19, 20, 29, 37, 49] have explored online approaches with the input of onboard sensor data. In [18, 22, 24, 35, 42], HDMap construction is treated as a segmentation task, relying on post-processing to achieve map vectorization. Later, Liao *et al.* [25, 26] represent map elements as ordered points and employ transformers [41, 57] after the BEV encoder for auto-regressive detection. Other modeling methods, such as Bezier curves and pivot points, are explored in BeMapNet [37] and PivotNet [6] for geometric representation. Several works [4, 19, 40, 43, 50] are proposed for different temporal fusion approaches. GeMap [54] and MapQR [31] improve mapping performance from the perspectives of geometric relationships and decoder design, respectively. In [5, 28, 51, 52], rasterized maps are utilized for structure compensation. Despite improved performance, these approaches seldom consider the generalization capability of models in unfamiliar driving scenarios. In this work, we incorporate uncertainty-instructed PV detection retracing, aiming for better performance across diverse scenarios.

PV Detection for 3D Perception. Instances from PV detection can be effective compensations for 3D perceptions [12, 14, 32, 47, 48, 53]. For 3D object detection, BEV-FormerV2 [48] utilizes a two-stage BEV detector that integrates the query from the PV head into the BEV space. In SimMoD [53], 2D detection is utilized to construct proposals for 3D fine-grained perception. Far3D [14] combines 2D detection boxes with estimated depth information to improve long-range perception performance. For model generalization, Lu *et al.* [32] propose a generalizable framework based on perspective debiasing, in which PV consistency learning is employed among different domains. Applying for 3D lanes, Li *et al.* [17] use 2D priors for lane detection and topology reasoning. Inspired by the above methods, we employ PV structural features for more generalizable HD map vectorization, in which distillations from 2D instances are utilized to ensure inference efficiency.

Uncertainty in Model. Generally, uncertainty refers to the model’s confidence level for its predictions, which can be a statistical representation to indicate the model’s under-

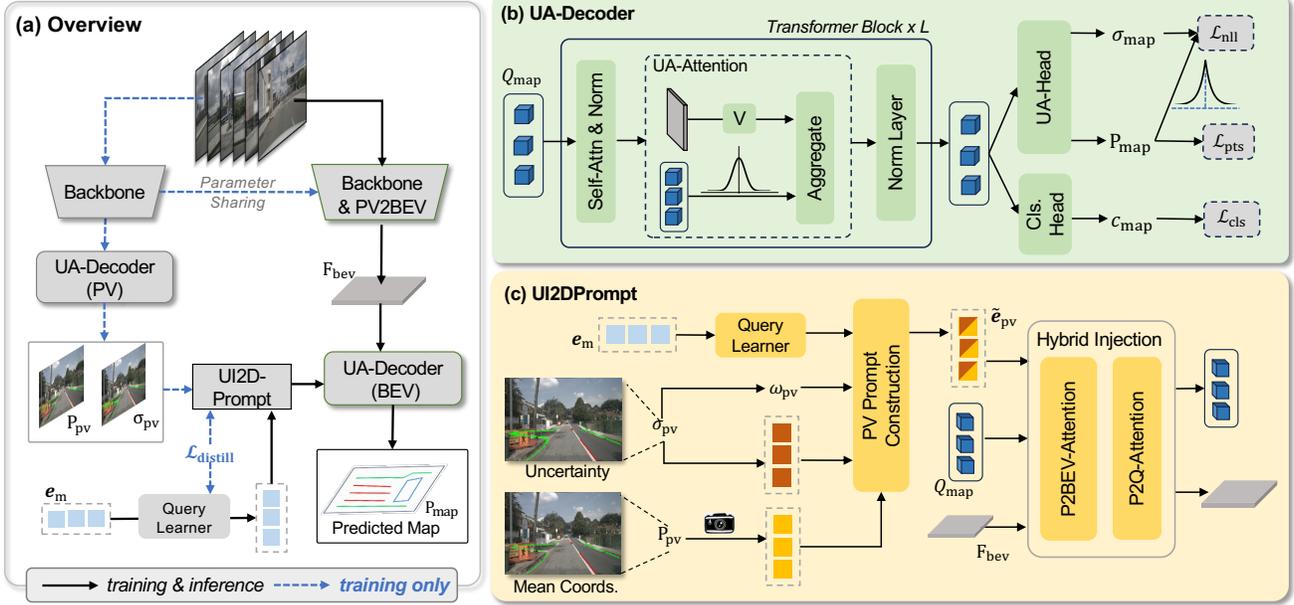


Figure 2. (a) **Overview of Our UIGenMap.** For training, the PV branch is introduced with the uncertainty-instructed structural injection, in which the MQ-Distillation is designed to mimic PV structural features. (b) **Uncertainty-Aware Decoder Architecture.** For each layer, UA-Decoder comprises probabilistic UA-Attention and UA-Head for reliable output. (c) **UI2DPrompt Design.** We construct PV prompts from PV-detected elements and their corresponding uncertainties, which are integrated with the main branch by hybrid injection.

standing of scenes [3, 15, 44]. Uncertainty has gained intensive attention due to the potential for reliability evaluation and performance enhancement [2, 16]. Specifically, the generalization problem with uncertainty modeling is investigated in [21, 36], which expects uncertainty estimation for out-of-distribution generalization. For the transformer structure [8, 34], uncertainty has been explored to guide the training and inference processes for more robust model performance. In the field of autonomous driving, significant uncertainty has been introduced for 3D occupancy [2, 44], optical flow [33], and object tracking [55]. Recently, Gu *et al.* [7] present uncertainty estimation for the vectorized map, which is targeted for trajectory prediction. In our design, we utilize uncertainty for more generalized HD map construction, which considers the variability in feature distributions and incorporates explicit instance features.

3. UIGenMap

Given surrounding-view images captured from onboard cameras, we aim to construct more robust vectorized map elements in the BEV space. Each map element comprises a class label c_{map} and an ordered sequence of points $P_{\text{map}} = \{(p_x^i, p_y^i)\}_{i=1}^N$ that represents the structure of the instance. N denotes the total number of points for each map element.

The architecture of UIGenMap is presented in Figure 2. Based on the main BEV detection, we obtain PV instance from an introduced pre-trained PV branch, which is injected as structural priors for BEV perception. Firstly, we extract

image features for the PV and BEV branches separately. As in [22], BEV features F_{bev} are constructed from the interaction between PV features and learnable BEV queries. Then, we design the uncertainty-aware (UA-) decoder from inner-layer UA-Attention and point-level heads to achieve reliable probabilistic outputs. Based on the PV coordinates and uncertainty output, we propose the UI2DPrompt module to construct reliable PV prompts. These prompts are incorporated into BEV features and map instance queries via hybrid injection, where uncertainties are utilized as the basis for feature selection and enhancement. Finally, the UA-Decoder employs enhanced BEV features and instance queries for reliable map predictions. To ensure real-time inference, the lightweight MQ-Distillation module performs query distillation from learned PV prompts to imitate explicit structural features, aiming for efficient computation.

3.1. Uncertainty-Aware Decoder

We design UA-Decoder for both PV and BEV space. Systematic uncertainty-aware design is employed to improve the statistical understanding of features and outputs to achieve more dynamic self-adjustment for different driving scenarios. To ensure comprehensive perception, each decoder layer comprises UA-Attention at the instance feature level and UA-Head for point-level uncertainty output along with positions, respectively.

UA-Attention. In conventional deformable attention [57], the weight α_i for each sampled feature values from BEV feature F_{BEV} is determined by linear layer ($\text{Linear}(\cdot)$) on

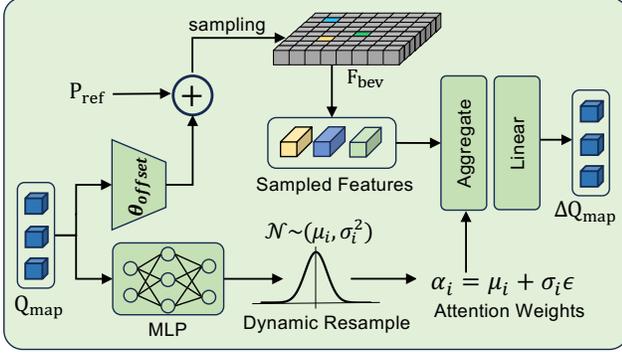


Figure 3. UA-Attention Design. The learned probabilistic weights α_i are dynamically resampled from Gaussian distribution with stochastic disturbance, which is multiplied with sampled feature values to update map queries.

Q_{map}^i , which can be expressed as $\alpha_i = \text{Linear}(Q_{\text{map}}^i)$, where Q_{map}^i is the i -th instance query of totally M map elements. Then, the updated part of query ΔQ_{map}^i is aggregated by the dot-product between the learned weight α_i and sampled feature values from F_{bev} . Considering complex driving scenarios, such deterministic weights are inadaptible for challenging situations. More dynamic weights are required to avoid potential ambiguities. Inspired by [8, 34], we explicitly model the attention scores under a Gaussian distribution $\alpha_i \sim \mathcal{N}(\mu_i, \sigma_i^2)$, as shown in Figure 3. We adopt the re-parameterization strategy to perform the probabilistic attention, which is constructed by the corresponding mean μ_i and variance σ_i . Such probabilistic parameters are predicted through multilayer perceptron ($\text{MLP}(\cdot)$). Thus, the construction of dynamic resampled attention weights can be formulated as below:

$$\{\mu_i, \sigma_i^2\} = \text{MLP}(Q_{\text{map}}^i), \quad (1)$$

$$\alpha_i = \mu_i + \sigma_i \epsilon, \epsilon \sim \mathcal{N}(0, 1), \quad (2)$$

where ϵ is the Gaussian noise conforming to normal distribution. During inference, attention weight α_i can be dynamically resampled from learned statistical distribution to suit for driving scenarios with large variations. Then query features can be obtained by dot production with the sampled feature values in order to enable reliable query updates.

UA-Head with Uncertainty Loss. Based on the robust design of UA-Attention, quantized uncertainty for each lane point output is also essential, which is point-specific at a fine-grained level to explicitly reflect the reliability and uncertainty information within the internal transformer layers. Inspired by [7], we replace the vanilla deterministic output with the UA-Head that predicts coordinates $\mathbf{P}_{\text{map}} = \{(\hat{p}_x^i, \hat{p}_y^i)\}_{i=1}^N$ along with their corresponding uncertainties $\sigma_{\text{map}} = \{(\sigma_x^i, \sigma_y^i)\}_{i=1}^N$ for each map element, as shown in Figure 2 (b). These parameters represent the location and scale of the joint Laplace distribution $\text{Laplace}(\hat{p}_j^i, \frac{\sigma_j^i}{\sqrt{2}})$ at

the j -th dimension (x or y) of the i -th point, which refers to a probabilistic representation for stable output.

For uncertainty-aware training, we combine probabilistic Negative Log-Likelihood (NLL) loss \mathcal{L}_{nll} with point regression loss \mathcal{L}_{pts} in [50] to balance uncertainties and coordinate outputs, in which \mathcal{L}_{nll} can be formulated as follows:

$$\mathcal{L}_{\text{nll}} = \lambda_{\text{nll}} \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^2 \left(\log(2\sigma_j^i) + \frac{|p_j^i - \hat{p}_j^i|}{\sigma_j^i} \right), \quad (3)$$

with λ_{nll} as the loss weight setting.

As illustrated in Figure 2 (a), UA-Decoder is used for the construction of the PV branch and the BEV map. To provide more robust structural features in PV branch, PV coordinates $\mathbf{P}_{\text{pv}} = \{x_{\text{pv}}^i, y_{\text{pv}}^i\}_{i=1}^N$ and the uncertainty parameters $\sigma_{\text{pv}} = \{\sigma_{\text{pv},x}^i, \sigma_{\text{pv},y}^i\}_{i=1}^N$ are constructed by the UA-Decoder structure, which is later utilized in our UI2DPrompt design. Thus, the systematic combination of uncertainty-aware design provides dynamic sampling during the learning process, along with established uncertainty, to construct a reliable basis for confidence scoring.

3.2. UI2DPrompt

PV branches contain abundant structural and geometric details. To mitigate information loss and misinterpretation during query-based PV-to-BEV transformation, reliable structural PV instances can be injected as compensation for BEV prediction. This section focuses on constructing PV prompts based on uncertainty-aware PV predictions, which are dynamically integrated with the map decoder for BEV map element detection.

Prompt Construction. Given the learned PV instance coordinates and the corresponding uncertainty parameters $\{\mathbf{P}_{\text{pv}}, \sigma_{\text{pv}}\}$, PV prompts are constructed to provide dynamic 2D priors. Firstly, we employ the estimated classification score c_{map} for PV candidates selection. When $c_{\text{map}} > c_{\text{thr}}$, PV instances are chosen as candidates for reliable localized information. Considering the sparse depth information and the characteristics of static map elements on the ground, we utilize Inverse Perspective Mapping (IPM) to transform \mathbf{P}_{pv} from the PV image coordinate to the BEV coordinate system $\mathbf{P}_{\text{trans}} = \{p'_x, p'_y\}$, which makes a flat grounding assumption and utilizes the affine transformation between the PV and BEV coordinates. Given the corresponding intrinsic parameters \mathbf{K}_{cam} and extrinsic transformation matrix $\mathbf{T}_{\text{ego2cam}}$, the projection \mathbf{P}_{pv} can be defined as:

$$\mathbf{P}_{\text{trans}} = \text{Proj}_{2\text{D}} \left(\mathbf{T}_{\text{ego2cam}}^{-1} \cdot \begin{bmatrix} \mathbf{K}_{\text{cam}}^{-1} \cdot \mathbf{P}_{\text{pv}} \\ 1 \end{bmatrix} \right), \quad (4)$$

where $\text{Proj}_{2\text{D}}(\cdot)$ denotes the conversion from homogeneous coordinates into 2D plane coordinates.

To ensure reliable enhanced features, the corresponding uncertainty parameters σ_{pv} cooperate as weights and feature

embeddings. We encode the transformed coordinates $\mathbf{P}_{\text{trans}}$ and uncertainty parameters σ_{pv} by $\phi_p(\cdot)$ and $\phi_\sigma(\cdot)$, respectively. Point-specific embeddings $\{\tilde{e}_{\text{pv}}^i\}_{i=1}^N$ are constructed by concatenation:

$$\tilde{e}_{\text{pv}}^i = \text{Concat}([\phi_p(\mathbf{P}_{\text{trans}}), \phi_\sigma(\sigma_{\text{pv}})]). \quad (5)$$

Thus, the structural PV output coordinates and uncertainty information are dynamically encoded.

Based on the above, PV uncertainties can serve as the quantized weights for dynamic feature strengthening and adjustment. Reliable confidence indicator is able to be formulated as the weight $\{\omega_{\text{pv}}^i\}_{i=1}^N$ concerning each point-specific prompt. Then, the enhanced PV prompt \tilde{e}_{pv}^i can be constructed by:

$$\omega_{\text{pv}}^i = \exp\left(\frac{(\|\sigma_{\text{pv}}^i\|_2)^{-1}}{\sum_{i=1}^N (\|\sigma_{\text{pv}}^i\|_2)^{-1}}\right), \quad \tilde{e}_{\text{pv}}^i = \omega_{\text{pv}}^i \tilde{e}_{\text{pv}}^i + e_{\text{m}}^i, \quad (6)$$

where $\exp(\cdot)$ represents the exponential amplification of the normalized inverse uncertainty value. $\{e_{\text{m}}^i\}_{i=1}^N$ is the predefined mimic query, which is utilized for MQ-Distillation. This approach reduces the introduction of PV errors, ensuring that injecting PV instances amplifies reliable PV information while minimizing uncertainty-induced errors.

Hybrid Injection. Constructed from PV instances with uncertainty output, PV prompts $\{\tilde{e}_{\text{pv}}^i\}_{i=1}^N$ are incorporated as compensation for the construction of the BEV map. These prompts are injected into the BEV feature F_{bev} and the map queries Q_{map}^i for the map decoder by the designed hybrid injection mechanism, which includes both prompt-to-BEV (P2BEV) and prompt-to-query (P2Q) attention.

As shown in Figure 4, point-level PV prompts $\{\tilde{e}_{\text{pv}}^i\}_{i=1}^N$ are integrated with F_{bev} through cross-attention of P2BEV. This step enriches BEV features by integrating detailed, point-specific information from PV prompts, enhancing instance-specific detail in the BEV representation:

$$F'_{\text{bev}} = \text{Attn}(W_q F_{\text{bev}}, W_k \tilde{e}_{\text{pv}}^i, W_v \tilde{e}_{\text{pv}}^i), \quad (7)$$

where $W_{\{q,k,v\}}$ are the corresponding projection matrix. Then, we encode point-level prompts into instance-level representations through a linear block $\phi_e(\cdot)$, which aggregates point-specific information to create a higher-level representation for subsequent processing. These instance-level prompts \tilde{e}_{inst} are used to integrate with Q_{map} , which inject reliable priors for the initialized queries and enhance the overall map construction:

$$\tilde{e}_{\text{inst}} = \phi_e(\{\tilde{e}_{\text{pv}}^i\}_{i=1}^N), \quad (8)$$

$$Q'_{\text{map}} = \text{Attn}(W_q Q_{\text{map}}, W_k \tilde{e}_{\text{inst}}, W_v \tilde{e}_{\text{inst}}). \quad (9)$$

By injecting PV priors into both BEV and decoder queries, this hybrid mechanism improves the reliability of BEV perception and reduces transition errors. Explicit structural

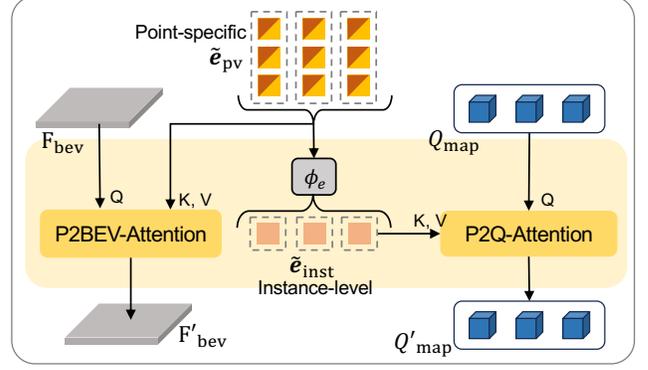


Figure 4. Design of Hybrid Injection. P2BEV-Attention and P2Q-Attention are designed for PV prompts to integrate with BEV features and map instance queries as structural compensation.

features enhance the generalization capability of the learned model and perform robustly in various driving environments, ensuring better accuracy and consistency.

3.3. Mimic Query Distillation

To address the extra computation introduced by the PV branch, we propose the lightweight MQ-Distillation to distill from PV prompt construction, where the mimic query can serve as an alternative to achieve real-time inference.

In MQ-Distillation, we define a set of mimic queries $\{e_{\text{m}}^i\}_{i=1}^N$, along with an MLP-based query learner $h(\cdot)$, which is designed to capture and distill the structural features of the PV prompt construction process \tilde{e}_{pv}^i . As shown in Equation 6, during training, the predefined mimic queries e_{m}^i and the constructed PV prompts \tilde{e}_{pv}^i are jointly combined by addition, which are sent to the hybrid injection for further operations.

To obtain sufficient structural features and achieve progressive alignment with PV prompts, we design the loss function that guides mimic queries for distillation. We use Mean Squared Error (MSE) as the distillation loss to measure the difference between the constructed PV prompts and the mimicked queries, which can be formulated as follows:

$$\mathcal{L}_{\text{distill}} = \lambda_{\text{distill}} \frac{1}{N} \sum_{i=1}^N (\tilde{e}_{\text{pv}}^i - h(e_{\text{m}}^i))^2, \quad (10)$$

where λ_{distill} is a weight coefficient that balances the impact of the distillation loss on the overall training process. Such loss function guides the mimic queries to accurately replicate the PV prompts generated by the PV branch, making them reliable substitutes for inference.

To this end, only the mimic queries can be utilized during inference, which eliminates the requirement for PV prompt computation. This approach maintains real-time inference capability by efficiently mimicking PV structural features with minimal computational cost, enabling streamlined yet robust map construction.

Method	Backbone	Epochs	Region-Based				City-Based				FPS
			AP _{ped}	AP _{div}	AP _{bou}	mAP	AP _{ped}	AP _{div}	AP _{bou}	mAP	
VectorMapNet* [ICML'23]	R50	120	15.8	17.0	21.2	18.0	-	-	-	-	-
MapTR* [ICLR'23]	R50	24	6.4	20.7	35.5	20.9	11.1	12.9	21.1	15.0	15.8
MGMap [CVPR'24]	R50	24	9.4	27.0	38.1	24.8	10.9	17.8	24.7	17.8	11.5
MapQR [ECCV'24]	R50	24	8.4	25.7	39.9	24.7	10.2	15.8	25.5	17.2	12.1
GeMap [ECCV'24]	R50	24	11.8	26.6	43.6	27.3	12.9	16.3	26.7	18.6	11.6
MapTRv2* [IJCV'24]	R50	24	14.7	28.7	43.3	28.9	18.9	18.6	27.9	21.8	12.9
StreamMapNet* [WACV'24]	R50	24	<u>32.2</u>	29.3	40.8	34.1	18.7	16.4	22.7	19.3	13.3
<i>UIGenMap-d (Ours)</i>	R50	24	40.3	30.8	46.8	39.3 (+5.2)	23.1	19.6	25.4	22.7 (+3.4)	12.2
<i>UIGenMap (Ours)</i>	R50	24	40.3	31.8	47.2	39.8 (+5.7)	24.1	20.7	26.1	23.6 (+4.3)	8.2
StreamMapNet [WACV'24]	SwinT	24	32.5	27.4	41.7	33.9	21.8	19.7	25.8	22.4	9.3
<i>UIGenMap-d (Ours)</i>	SwinT	24	42.6	31.9	47.2	40.6 (+6.7)	24.7	21.2	28.8	24.9 (+2.5)	8.5

Table 1. Quantitative comparisons on *nuScenes* validation dataset, utilizing region-based and city-based dataset partitions. Methods with “*” denote that data results are obtained from corresponding papers [27, 50], and others are reproduced by official public codes. “R50” and “SwinT” denotes the image backbone of ResNet50 [9] and Swin-Tiny [30]. “*UIGenMap-d*” represents the utilization of mimic query from MQ-Distillation. FPS is measured on a single RTX A40 GPU.

3.4. Training and Inference

During the training process, combined with \mathcal{L}_{nl} for uncertainty output and $\mathcal{L}_{\text{distill}}$ for mimic query distillation, the overall loss setting \mathcal{L}_{map} can be represented as:

$$\mathcal{L}_{\text{map}} = \lambda_1 \mathcal{L}_{\text{pts}} + \lambda_2 \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{nl}} + \mathcal{L}_{\text{distill}}, \quad (11)$$

where \mathcal{L}_{pts} denotes the point-to-point Manhattan distance between the regressed points and the ground-truth. \mathcal{L}_{cls} represents the focal loss for map classification. λ_1 and λ_2 are the corresponding loss weights for point regression and map classification, respectively.

For inference, the well-learned mimic query can only be employed as a substitute for PV prompts, avoiding the computational burden for the PV detection branch. Thus, our approach ensures real-time inference without introducing unnecessary computational overhead, while preserving effective performance. Additionally, the model’s uncertainty learning enables dynamic sampling during inference, enhancing the adaptability to diverse driving environments.

4. Experiments

4.1. Datasets and Metrics

Datasets. We conduct extensive experiments on widely used autonomous driving datasets, including *nuScenes* [1] and *Argoverse2* [46]. *nuScenes* dataset comprises 1,000 driving scenes from Boston and Singapore. Each scene lasts approximately 20 seconds with a 2Hz sampling rate. For each sample, there are six surrounding images and 32-beam LiDAR point clouds. *Argoverse2* contains 1,000 scenes from 6 cities, with each scene capturing 15 seconds of data at 20Hz from 7 cameras and 10Hz LiDAR sweeps. In our experiments, the sampling frequency for *Argoverse2* is unified into 2Hz as in *StreamMapNet* [50].

Geo-based Data Partitions. Considering the overlap of widely used training and validation partitions, we adopt

geographically disjoint (geo-based) dataset partitions as benchmarks to evaluate the generalization capability of the presented model. Experiments are conducted on both **region-based** and **city-based** data partitions. The region-based partition is adopted from [50], where the driving locations geographically disjoin the training and validation sets. In the city-based data partition, we apply city-wise splits in [27], where the training and validation data are from different cities with a larger distribution change. On *nuScenes*, Boston and Onenorth are used for training, while Queenstown and Holland Village are used for validation. For *Argoverse2*, Miami and Pittsburgh are used in training, and the remains are employed for validation. For different dataset partitions, models are separately trained on that split’s training set and evaluated on their corresponding validation set.

Evaluation Metrics. In line with the existing works [26, 50], we conduct experiments on three static map elements, namely *lane dividers*, *pedestrian crossings*, and *road boundaries*. Perception range is set to $([-15.0m, 15.0m])$ for X-axis and $([-30.0m, 30.0m])$ for Y-axis. We use average precision as the evaluation metric, which calculates the Chamfer distance between predictions and ground truths. The mean Average Precision (mAP) is obtained by averaging AP across the distinct thresholds: $\{0.5m, 1m, 1.5m\}$.

4.2. Implementation Details

For fair comparisons, we employ ResNet50 [9] as the backbone. As in [50], the size of BEV features is set to 50×100 . The number of map queries and query dimensions are set to 100 and 512. We use AdamW optimizer with a learning rate of $4e^{-4}$. Training epochs are limited to prevent potential overfitting. All models are trained on 4 NVIDIA Tesla V100 GPUs with a batch size of 4. The weighted factors λ_1 and λ_2 are set to 50.0 and 5.0, separately. λ_{nl} and λ_{distill} are set to 0.05 and 10.0. We adopt a pre-trained PV detection branch for our *UIGenMap* training process. More details

Method	Backbone	Epochs	Region-Based				City-Based				FPS
			AP _{ped}	AP _{div}	AP _{bou}	mAP	AP _{ped}	AP _{div}	AP _{bou}	mAP	
VectorMapNet* [ICML'23]	R50	120	35.6	34.9	37.8	36.1	-	-	-	-	-
MapTR* [ICLR'23]	R50	30	48.1	50.4	55.0	51.1	34.7	41.7	34.7	37.9	15.8
GeMap [ECCV'24]	R50	30	52.5	59.2	58.2	56.6	34.5	41.1	40.4	38.7	11.6
MapTRv2* [IJCV'24]	R50	30	53.3	58.7	58.8	56.9	37.4	42.2	41.9	40.5	12.9
MapQR [ECCV'24]	R50	30	54.2	58.0	60.4	57.5	33.8	41.5	41.6	38.9	12.1
StreamMapNet* [WACV'24]	R50	30	57.9	55.7	61.3	58.3	40.2	43.4	43.3	42.3	13.3
<i>UIGenMap-d (Ours)</i>	R50	30	<u>58.3</u>	58.4	<u>62.4</u>	<u>59.7 (+1.4)</u>	<u>42.9</u>	<u>44.4</u>	<u>44.7</u>	<u>44.0 (+1.7)</u>	12.2
<i>UIGenMap (Ours)</i>	R50	30	59.8	<u>58.7</u>	62.7	60.4 (+2.1)	43.7	45.7	45.1	44.8 (+2.5)	8.2

Table 2. Quantitative comparisons on *Argoverse2* validation datasets, utilizing region-based and city-based dataset partitions. Methods with “*” denote that data results are obtained from corresponding papers [27, 50], and others are reproduced by official public codes. Models are trained for 30 epochs and follow the same setting. FPS is measured on a single RTX A40 GPU.

Dataset	Method	Epochs	AP _{ped}	AP _{div}	AP _{bou}	mAP
nuScenes	MapTRV2 [26]	24	59.8	62.4	62.4	61.5
	StreamMapNet [50]	24	60.4	61.9	58.9	60.4
	SQD-MapNet [43]	24	63.0	65.5	63.3	63.9
	<i>UIGenMap(Ours)</i>	24	64.2	65.4	64.9	64.8
Argoverse2	StreamMapNet [50]	30	62.9	59.5	63.0	61.5
	SQD-MapNet [43]	30	64.9	60.2	64.9	63.3
	<i>UIGenMap(Ours)</i>	30	64.5	60.8	66.1	63.8

Table 3. Performance on nuScenes and Argoverse2 datasets with the original dataset partition at 30m range.

bkb	UI2DPrompt	UA-Decoder	AP _{ped}	AP _{div}	AP _{bou}	mAP
			32.2	29.3	40.8	34.1
	✓		39.9	28.6	43.8	37.4
		✓	36.1	31.3	44.0	37.1
	✓	✓	41.6	29.7	44.8	38.7
✓	✓		41.4	29.8	47.4	39.5
✓		✓	39.2	30.1	46.2	38.5
✓	✓	✓	40.3	31.8	47.2	39.8

Table 4. Ablations on the impact of component designs, ‘bkb’ means the shared image backbone from PV branches.

are provided in the *supplementary materials*.

4.3. Main Results

Results on nuScenes. We compare the generalization capability with previous methods [25, 26, 28, 31, 50, 54] on geospatial data partitions, in which the image size is set to 480×800. As shown in Table 1, our UIGenMap outperforms previous methods at a large margin. For the ResNet50 backbone, UIGenMap achieves 5.7 mAP improvement on the region-based partition, and 4.3 mAP gain is obtained on the city-based splits. For more efficient design, UIGenMap-d also archives 39.3 mAP with a faster inference speed. Experiments on the Swin-T [30] backbone show significant performance with 40.6 mAP achievement. Furthermore, we show some visual examples in Figure 5, indicating more intuitive achievements for UIGenMap.

Results on Argoverse2. More experiments are conducted on Argoverse2 dataset. The image size is 608×608, and the sampling rate is 2Hz. For a fair comparison, the models are trained for 30 epochs. As shown in Table 2, UIGen-

Map achieves the best results of 60.4 mAP on region-based data partitions and 43.4 mAP on city-based splits, which performs better than previous approaches. Upon further investigation, calibration errors and dataset distribution also affect model performance. More visualization can be found in our *supplementary materials*.

Results on Original Data Partitions. For thorough evaluation, we conducted experiments on the original data partitions with similar driving scenarios, which are widely used in previous studies. As shown in Table 3, we compare our methods with others utilizing the same baseline. UIGenMap still outperforms the existing approaches, which outperforms StreamMapNet with +4.4 mAP and +2.3 mAP separately, validating the effectiveness of our approach.

4.4. Ablation Studies

In this section, ablations are conducted to verify the effectiveness of our designs. We perform experiments on nuScenes with region-based partition using ResNet50 backbone. All models are trained for 24 epochs. More ablations are provided in the *supplementary materials*.

Contributions of Main Components. Table 4 demonstrates the impact of each component in our design. Based on the pre-trained PV backbone, we design the UA-Decoder for more reliable uncertainty output and UI2DPrompt to dynamically inject PV structural information, which achieves +2.7 mAP and +2.4 mAP improvement, respectively. The combination of these designs gains +4.0 mAP total. The ablations are conducted on UA-Decoder at BEV space. Besides, the pre-trained PV backbone further emphasizes the effectiveness to 39.8 mAP for better scene understanding.

Design of UA-Decoder. We design the UA-Attention for probabilistic attention and the UA-Head to construct the uncertainty output from the instance and finer point levels. Table 5 shows the ablation without other modules. Our uncertainty-aware design supports dynamic adaptation to achieve more robust performance, which achieves +3.0 mAP improvements based on systematic combination.

Design for UI2DPrompt. UI2DPrompt consists of prompt construction and a hybrid injection module. Table 6 shows

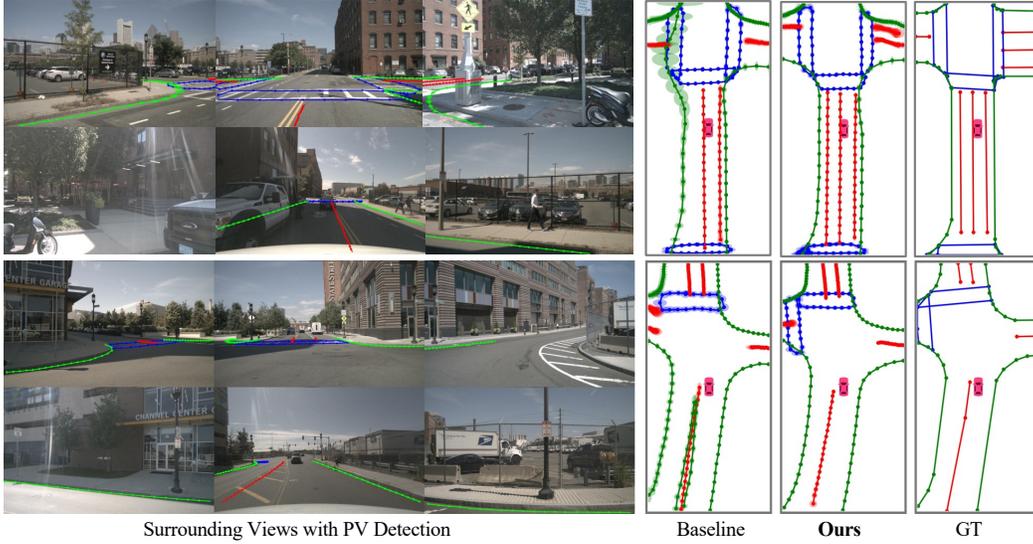


Figure 5. Qualitative results on the nuScenes dataset with the region-based data partition. Uncertainty outputs are represented as circles.

UA-Att.	UA-Head	AP _{ped}	AP _{div}	AP _{bou}	mAP
		32.2	29.3	40.8	34.1
✓		31.7	30.4	42.9	35.0
	✓	34.8	30.5	43.4	36.3
✓	✓	36.1	31.3	44.0	37.1

Table 5. Ablations on UA-Decoder without other modules. UA-Att and UA-Head represent the utilization of uncertainty-aware attention and output head respectively.

Coords. Trans.	Unc. Utilization	mAP
/	/	37.6
IPM	embedding only	39.3
IPM	weight & embedding	39.8
3D Trans.	weight & embedding	38.7

Table 6. Ablations on the construction of PV prompts. “3D Trans.” means the depth-based transformation. The first row with “/” indicates direct PV coordinate embedding.

the strategy selection for constructing PV prompts. Compared to a depth-based 3D transition, IPM is more suitable for coordinating the transformation based on the assumption of a flattened road surface. When uncertainty is applied both as dynamic weighting and as learning-based embeddings, performance is further enhanced by reliable PV prompts, achieving an improvement of +2.2 mAP.

Table 7 demonstrates the impact of the hybrid injection design, with a total performance of 39.8 mAP for combining prompt-to-BEV and prompt-to-query attention. This result confirms the optimal utilization of reliable structural information, effectively compensating for information loss and errors introduced in the basic learning process.

Ablations on MQ-Distillation. To verify the effectiveness of the distillation module, we conduct ablations on the hyperparameter of distillation loss weight λ_{distill} , as this affects the imitation learning of PV structural features. As

P2BEV	P2Q	AP _{ped}	AP _{div}	AP _{bou}	mAP
		39.3	30.4	45.5	38.4
✓		39.9	31.1	46.0	39.2
	✓	39.6	30.7	46.1	38.8
✓	✓	40.3	31.8	47.2	39.8

Table 7. Ablations on hybrid injection. P2BEV and P2Q mean the cross-attention for F_{bev} and Q_{map} separately.

λ_{distill}	5	7	10	15	/
mAP	37.5	38.2	39.3	38.8	39.8

Table 8. Ablations on λ_{distill} for MQ-Distillation Design. The last column refers to the use of the best teacher of PV prompts.

illustrated in Table 8, the best performance is achieved by setting the weight to 10.0, which refers to the best learning process of the mimic query for explicit structural features.

5. Conclusion and Future Work

In this paper, we propose UGenMap, an uncertainty-instructed structure injection approach to achieve a more generalized HD map construction. Specifically, we introduce the PV detection branch to obtain explicit structural features, in which the UA-Decoder produces adaptive probabilistic outputs. The UI2DPrompt module is proposed to construct reliable PV prompts and achieve dynamic integration through the hybrid injection designed. For real-time inference, MQ-Distillation is designed to mimic the structural prompts of PV. Comprehensive experiments on the challenging geo-based partitions of nuScenes and Argoverse2 datasets demonstrate the generalization of our approach.

For limitations, current methods are designed and evaluated primarily on specific datasets. In the future, developing larger and more diverse datasets and benchmarks across various locations remains a crucial area for exploration.

Acknowledgments

This work is supported by National Natural Science Foundation of China under Grant No.62376244. It is also supported by Information Technology Center and State Key Lab of CAD&CG, Zhejiang University.

References

- [1] 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 multimodal dataset for autonomous driving. In *CVPR*, pages 11621–11631, 2020. 1, 2, 6
- [2] Anh-Quan Cao, Angela Dai, and Raoul de Charette. Pasco: Urban 3d panoptic scene completion with uncertainty awareness. In *CVPR*, pages 14554–14564, 2024. 3
- [3] Jie Chang, Zhonghao Lan, Changmao Cheng, and Yichen Wei. Data uncertainty learning in face recognition. In *CVPR*, pages 5710–5719, 2020. 3
- [4] Jiacheng Chen, Yuefan Wu, Jiaqi Tan, Hang Ma, and Yasutaka Furukawa. Maptracker: Tracking with strided memory fusion for consistent vector hd mapping. In *ECCV*, pages 90–107. Springer, 2025. 1, 2
- [5] Sehwan Choi, Jungho Kim, Hongjae Shin, and Jun Won Choi. Mask2map: Vectorized hd map construction using bird’s eye view segmentation masks. In *ECCV*, 2024. 2
- [6] Wenjie Ding, Limeng Qiao, Xi Qiu, and Chi Zhang. Pivotnet: Vectorized pivot learning for end-to-end hd map construction. In *ICCV*, pages 3672–3682, 2023. 2
- [7] Xunjiang Gu, Guanyu Song, Igor Gilitschenski, Marco Pavone, and Boris Ivanovic. Producing and leveraging online map uncertainty in trajectory prediction. In *CVPR*, pages 14521–14530, 2024. 3, 4
- [8] Hongji Guo, Hanjing Wang, and Qiang Ji. Uncertainty-guided probabilistic transformer for complex action recognition. In *CVPR*, pages 20052–20061, 2022. 3, 4
- [9] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 6
- [10] Yuzhe He, Shuang Liang, Xiaofei Rui, Chengying Cai, and Guowei Wan. Egovm: Achieving precise ego-localization using lightweight vectorized maps. *arXiv preprint arXiv:2307.08991*, 2023. 1
- [11] Haotian Hu, Fanyi Wang, Yaonong Wang, Laifeng Hu, Jingwei Xu, and Zhiwang Zhang. Admap: Anti-disturbance framework for reconstructing online vectorized hd map. *arXiv preprint arXiv:2401.13172*, 2024. 2
- [12] Haoxuanye Ji, Pengpeng Liang, and Erkang Cheng. Enhancing 3d object detection with 2d detection-guided query anchors. In *CVPR*, pages 21178–21187, 2024. 2
- [13] Bo Jiang, Shaoyu Chen, Qing Xu, Bencheng Liao, Jiajie Chen, Helong Zhou, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Vad: Vectorized scene representation for efficient autonomous driving. In *ICCV*, pages 8340–8350, 2023. 1
- [14] Xiaohui Jiang, Shuailin Li, Yingfei Liu, Shihao Wang, Fan Jia, Tiancai Wang, Lijin Han, and Xiangyu Zhang. Far3d: Expanding the horizon for surround-view 3d object detection. In *AAAI*, pages 2561–2569, 2024. 2
- [15] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *NeurIPS*, 2017. 3
- [16] Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. In *NeurIPS*, 2017. 3
- [17] Han Li, Zehao Huang, Zitian Wang, Wenge Rong, Naiyan Wang, and Si Liu. Enhancing 3d lane detection and topology reasoning with 2d lane priors. *arXiv preprint arXiv:2406.03105*, 2024. 2
- [18] Qi Li, Yue Wang, Yilun Wang, and Hang Zhao. Hdmapnet: An online hd map construction and evaluation framework. In *ICRA*, pages 4628–4634. IEEE, 2022. 2
- [19] Siyu Li, Jiacheng Lin, Hao Shi, Jiaming Zhang, Song Wang, You Yao, Zhiyong Li, and Kailun Yang. Dtlcmapper: Dual temporal consistent learning for vectorized hd map construction. *arXiv preprint arXiv:2405.05518*, 2024. 2
- [20] Siyu Li, Kailun Yang, Hao Shi, Song Wang, You Yao, and Zhiyong Li. Genmapping: Unleashing the potential of inverse perspective mapping for robust online hd map construction. *arXiv preprint arXiv:2409.08688*, 2024. 2
- [21] Xiaotong Li, Yongxing Dai, Yixiao Ge, Jun Liu, Ying Shan, and Lingyu Duan. Uncertainty modeling for out-of-distribution generalization. In *ICLR*, 2022. 3
- [22] 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 *ECCV*, pages 1–18. Springer, 2022. 2, 3
- [23] Zhiqi Li, Zhiding Yu, Shiyi Lan, Jiahan Li, Jan Kautz, Tong Lu, and Jose M Alvarez. Is ego status all you need for open-loop end-to-end autonomous driving? In *ICCV*, pages 14864–14873, 2024. 1
- [24] Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. In *NeurIPS*, pages 10421–10434, 2022. 2
- [25] Bencheng Liao, Shaoyu Chen, Xinggang Wang, Tianheng Cheng, Qian Zhang, Wenyu Liu, and Chang Huang. Maptr: Structured modeling and learning for online vectorized hd map construction. In *ICLR*, 2022. 1, 2, 7
- [26] Bencheng Liao, Shaoyu Chen, Yunchi Zhang, Bo Jiang, Qian Zhang, Wenyu Liu, Chang Huang, and Xinggang Wang. Maptrv2: An end-to-end framework for online vectorized hd map construction. *IJCV*, pages 1–23, 2024. 1, 2, 6, 7
- [27] Adam Lilja, Junsheng Fu, Erik Stenborg, and Lars Hammarstrand. Localization is all you evaluate: Data leakage in online mapping datasets and how to fix it. In *CVPR*, pages 22150–22159, 2024. 1, 6, 7
- [28] Xiaolu Liu, Song Wang, Wentong Li, Ruizi Yang, Junbo Chen, and Jianke Zhu. Mgmmap: Mask-guided learning for online vectorized hd map construction. In *CVPR*, pages 14812–14821, 2024. 1, 2, 7
- [29] Yicheng Liu, Tianyuan Yuan, Yue Wang, Yilun Wang, and Hang Zhao. Vectormapnet: End-to-end vectorized hd map learning. In *ICML*, pages 22352–22369. PMLR, 2023. 2

- [30] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. 6, 7
- [31] Zihao Liu, Xiaoyu Zhang, Guangwei Liu, Ji Zhao, and Ningyi Xu. Leveraging enhanced queries of point sets for vectorized map construction. *arXiv preprint arXiv:2402.17430*, 2024. 1, 2, 7
- [32] Hao Lu, Yunpeng Zhang, Qing Lian, Dalong Du, and Yingcong Chen. Towards generalizable multi-camera 3d object detection via perspective debiasing. *arXiv preprint arXiv:2310.11346*, 2023. 2
- [33] Ao Luo, Fan Yang, Xin Li, Lang Nie, Chunyu Lin, Haoqiang Fan, and Shuaicheng Liu. Gafflow: Incorporating gaussian attention into optical flow. In *ICCV*, pages 9642–9651, 2023. 3
- [34] Jiahuan Pei, Cheng Wang, and György Szarvas. Transformer uncertainty estimation with hierarchical stochastic attention. In *AAAI*, pages 11147–11155, 2022. 3, 4
- [35] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, pages 194–210. Springer, 2020. 2
- [36] Fengchun Qiao and Xi Peng. Uncertainty-guided model generalization to unseen domains. In *CVPR*, pages 6790–6800, 2021. 3
- [37] Limeng Qiao, Wenjie Ding, Xi Qiu, and Chi Zhang. End-to-end vectorized hd-map construction with piecewise bezier curve. In *CVPR*, pages 13218–13228, 2023. 2
- [38] Tixiao Shan and Brendan Englot. Lego-loam: Lightweight and ground-optimized lidar odometry and mapping on variable terrain. In *IROS*, pages 4758–4765. IEEE, 2018. 2
- [39] Tixiao Shan, Brendan Englot, Drew Meyers, Wei Wang, Carlo Ratti, and Daniela Rus. Lio-sam: Tightly-coupled lidar inertial odometry via smoothing and mapping. In *IROS*, pages 5135–5142. IEEE, 2020. 2
- [40] Jingyu Song, Xudong Chen, Liupei Lu, Jie Li, and Katherine A Skinner. Memfusionmap: Working memory fusion for online vectorized hd map construction. *arXiv preprint arXiv:2409.18737*, 2024. 2
- [41] A Vaswani. Attention is all you need. In *NeurIPS*, 2017. 2
- [42] Song Wang, Wentong Li, Wenyu Liu, Xiaolu Liu, and Jianke Zhu. Lidar2map: In defense of lidar-based semantic map construction using online camera distillation. In *CVPR*, pages 5186–5195, 2023. 2
- [43] Shuo Wang, Fan Jia, Yingfei Liu, Yucheng Zhao, Zehui Chen, Tiancai Wang, Chi Zhang, Xiangyu Zhang, and Feng Zhao. Stream query denoising for vectorized hd map construction. In *ECCV*, 2024. 2, 7
- [44] Song Wang, Zhongdao Wang, Jiawei Yu, Wentong Li, Bailan Feng, Junbo Chen, and Jianke Zhu. Reliocc: Towards reliable semantic occupancy prediction via uncertainty learning. *arXiv preprint arXiv:2409.18026*, 2024. 3
- [45] Song Wang, Jiawei Yu, Wentong Li, Wenyu Liu, Xiaolu Liu, Junbo Chen, and Jianke Zhu. Not all voxels are equal: Hardness-aware semantic scene completion with self-distillation. In *CVPR*, pages 14792–14801, 2024. 1
- [46] Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandelwal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemodel Pontes, Deva Ramanan, Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception and forecasting. In *NeurIPS*, 2021. 1, 2, 6
- [47] Qiangqiang Wu, Yan Xia, Jia Wan, and Antoni B Chan. Boosting 3d single object tracking with 2d matching distillation and 3d pre-training. In *ECCV*, 2024. 2
- [48] Chenyu Yang, Yuntao Chen, Hao Tian, Chenxin Tao, Xizhou Zhu, Zhaoxiang Zhang, Gao Huang, Hongyang Li, Yu Qiao, Lewei Lu, et al. Bevformer v2: Adapting modern image backbones to bird’s-eye-view recognition via perspective supervision. In *CVPR*, pages 17830–17839, 2023. 2
- [49] Jing Yang, Minyue Jiang, Sen Yang, Xiao Tan, Yingying Li, Errui Ding, Hanli Wang, and Jingdong Wang. Mgmmapnet: Multi-granularity representation learning for end-to-end vectorized hd map construction. *arXiv preprint arXiv:2410.07733*, 2024. 2
- [50] Tianyuan Yuan, Yicheng Liu, Yue Wang, Yilun Wang, and Hang Zhao. Streammapnet: Streaming mapping network for vectorized online hd map construction. In *WACV*, pages 7356–7365, 2024. 1, 2, 4, 6, 7
- [51] Gongjie Zhang, Jiahao Lin, Shuang Wu, Zhipeng Luo, Yang Xue, Shijian Lu, Zuoguan Wang, et al. Online map vectorization for autonomous driving: A rasterization perspective. In *NeurIPS*, 2024. 2
- [52] Xiaoyu Zhang, Guangwei Liu, Zihao Liu, Ningyi Xu, Yunhui Liu, and Ji Zhao. Enhancing vectorized map perception with historical rasterized maps. In *ECCV*, 2024. 2
- [53] Yunpeng Zhang, Wenzhao Zheng, Zheng Zhu, Guan Huang, Jiwen Lu, and Jie Zhou. A simple baseline for multi-camera 3d object detection. In *AAAI*, pages 3507–3515, 2023. 2
- [54] Zhixin Zhang, Yiyuan Zhang, Xiaohan Ding, Fusheng Jin, and Xiangyu Yue. Online vectorized hd map construction using geometry. *arXiv preprint arXiv:2312.03341*, 2023. 2, 7
- [55] Lijun Zhou, Tao Tang, Pengkun Hao, Zihang He, Kalok Ho, Shuo Gu, Wenbo Hou, Zhihui Hao, Haiyang Sun, Kun Zhan, et al. Ua-track: Uncertainty-aware end-to-end 3d multi-object tracking. *arXiv preprint arXiv:2406.02147*, 2024. 3
- [56] Yi Zhou, Hui Zhang, Jiaqian Yu, Yifan Yang, Sangil Jung, Seung-In Park, and ByungIn Yoo. Himap: Hybrid representation learning for end-to-end vectorized hd map construction. In *CVPR*, pages 15396–15406, 2024. 1
- [57] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. 1, 2, 3
- [58] Yanjin Zhu, Xin Zheng, and Jianke Zhu. Mesh-loam: Real-time mesh-based lidar odometry and mapping. *TTV*, 2024. 2