

SEF-MAP: Subspace-Decomposed Expert Fusion for Robust Multimodal HD Map Prediction

Haoxiang Fu^{1,*}, Lingfeng Zhang^{2,3,*}, Hao Li^{7,*}, Ruibing Hu⁴, Zhengrong Li⁵
Guanjing Liu⁶, Zimu Tan⁷, Long Chen³, Hangjun Ye³, Xiaoshuai Hao^{3,✉}

¹National University of Singapore

²Shenzhen Tsinghua International Graduate School, Tsinghua University

³Xiaomi EV, ⁴Chinese University of Hong Kong

⁵The University of Manchester, ⁶Renmin University of China

⁷Independent Researcher

haoxiaoshuai714@163.com

Abstract—High-definition (HD) maps are essential for autonomous driving, yet multi-modal fusion often suffers from inconsistency between camera and LiDAR modalities, leading to performance degradation under low-light conditions, occlusions, or sparse point clouds. To address this, we propose *SEF-MAP*, a Subspace-Expert Fusion framework for robust multimodal HD map prediction. The key idea is to explicitly disentangle BEV features into four semantic subspaces: LiDAR-private, Image-private, Shared, and Interaction. Each subspace is assigned a dedicated expert, thereby preserving modality-specific cues while capturing cross-modal consensus. To adaptively combine expert outputs, we introduce an uncertainty-aware gating mechanism at the BEV-cell level, where unreliable experts are down-weighted based on predictive variance, complemented by a usage balance regularizer to prevent expert collapse. To enhance robustness in degraded conditions and promote role specialization, we further propose distribution-aware masking: during training, modality-drop scenarios are simulated using EMA-statistical surrogate features, and a specialization loss enforces distinct behaviors of private, shared, and interaction experts across complete and masked inputs. Experiments on nuScenes and Argoverse2 benchmarks demonstrate that *SEF-MAP* achieves state-of-the-art performance, surpassing prior methods by +4.2% and +4.8% in mAP, respectively. *SEF-MAP* provides a robust and effective solution for multi-modal HD map prediction under diverse and degraded conditions.

I. INTRODUCTION

High-definition (HD) maps are crucial for autonomous driving, providing precise semantic and geometric information for perception, planning, and navigation [1]–[7]. Unlike raw sensor data, HD maps offer structured road topology, lane geometry, and traffic elements, supporting reliable long-term autonomy [2], [8]. The rise of bird’s-eye-view (BEV) perception has enabled HD map construction directly from multi-view cameras and LiDAR. However, building robust and accurate maps remains challenging due to the heterogeneity of modalities, sensor noise, and environmental uncertainties [1], [9].

Current methods for multi-modal HD map prediction often combine image and LiDAR features in BEV space using

feature concatenation or attention-based aggregation. These approaches overlook two key factors: (i) the complementary but modality-specific nature of image and LiDAR data, and (ii) the dynamic reliability of these modalities over space and time [10]. For example, while camera features excel at capturing lane markings, they fail under poor lighting, and LiDAR offers stable geometry but is limited by sparsity and occlusion. Without modeling these discrepancies, conventional fusion methods can lead to unreliable predictions when one modality is degraded [1].

To tackle the challenges of multimodal fusion and robustness under degraded conditions, we propose *SEF-MAP*, a Subspace-Expert Fusion framework. The central idea is to capture the unique strengths of each modality while reducing redundancy and enhancing complementarity. To this end, we explicitly decompose BEV features into four subspaces: LiDAR-private, Image-private, Shared, and Interaction. Each subspace is assigned to a specialized expert, which allows the model to retain modality-specific information while also learning cross-modal consensus. To integrate the outputs of these experts, we design an uncertainty-aware gating mechanism that dynamically adjusts expert contributions based on their estimated reliability, enabling more confident experts to play a greater role in the final prediction. To further improve robustness under degraded or missing modalities, we introduce distribution-aware masking during training: when a modality is dropped, its features are replaced by realistic surrogate samples derived from EMA statistics. This encourages each expert to specialize in its designated role and equips the framework to handle partial or corrupted observations. Finally, to stabilize the learning process and avoid expert collapse, we incorporate specialization losses that explicitly enforce the separation of modality-private, shared, and interaction knowledge. Together, these components make *SEF-MAP* a principled and effective solution for robust multimodal HD map prediction.

Our main contributions are summarized as follows:

- We develop a subspace-decomposed fusion framework that explicitly separates multimodal BEV features into LiDAR-private, Image-private, Shared, and Interac-

* These authors contributed equally to this work.

✉ Corresponding Author.

tion streams, thereby mitigating semantic misalignment across modalities.

- We propose a distribution-aware masking strategy coupled with specialization losses to enforce expert roles and improve robustness under unimodal degradation, directly addressing the fragility of fused representations when one modality is unreliable.
- We design an uncertainty-aware gating mechanism with balance regularizers, enabling adaptive expert selection while preventing redundancy and expert collapse, which alleviates the issue of redundancy and suppression between modalities.
- We validate our framework on nuScenes and Argoverse2 and report consistent improvements over state-of-the-art baselines, with +4.2% mAP on nuScenes and +4.8% on Argoverse2.

II. RELATED WORK

A. HD Map Construction Techniques

HD maps provide crucial prior environmental information for autonomous driving systems. The core task of HD map construction is to predict a collection of vectorized and topology-preserving static map elements (e.g., lane dividers, road boundaries, pedestrian crossings) in BEV space [11]–[14]. HD map construction approaches can be categorized into three main types: **Camera-based methods**, **LiDAR-based methods**, and **Fusion-based methods**. Camera-based methods [15]–[18] leverage geometric priors to project image features into BEV space, often introducing spatial distortions. These methods typically rely on high-resolution imagery and large-scale pretrained models to improve accuracy [19]–[21], limiting their practical applicability [22], [23]. LiDAR-based methods [24]–[26] provide precise geometric measurements for unified BEV representations but are constrained by point sparsity and environmental interference, compromising perception reliability [10], [27]. Fusion-based methods attempt to combine the complementary strengths of both modalities through feature concatenation or attention-based aggregation [28]. However, existing fusion approaches often suffer from modality misalignment and fail to adequately handle scenarios where one modality is degraded or unreliable, leading to suboptimal performance under challenging conditions [27]. Our approach tackles these challenges by decomposing multimodal features into specialized subspaces that preserve modality-specific information while enabling robust cross-modal fusion under degraded conditions.

B. Multimodal Fusion Techniques

Multimodal fusion, which integrates heterogeneous sensor data such as cameras and LiDAR, has become central to autonomous driving perception [29], [30]. Early approaches such as MV3D and AVOD [31], [32] relied on direct feature concatenation. While simple, these methods failed to capture meaningful cross-modal correlations and degraded severely under sensor failures [33], [34]. To overcome these shortcomings, more advanced designs emerged. Transformer-based methods like TransFusion [35] employ soft associations

between modalities, enabling flexible alignment but with high computational cost. BEV-centric frameworks such as BEVFusion [36]–[38] project multi-modal features into a unified BEV space and adaptively combine them through gating networks, improving robustness but incurring significant data and compute demands. Lightweight approaches, exemplified by SparseFusion [39], attempt to balance accuracy and efficiency by leveraging sparse attention mechanisms.

In the context of HD map construction, recent methods have further emphasized temporal reasoning and cross-modal interactions. MemFusionMap [23] introduces memory modules to strengthen sequential modeling, while InteractionMap [40] enhances fine-grained feature exchanges across modalities. Despite their improvements, these approaches remain computationally intensive and still lack explicit modeling of heterogeneous modality-specific cues. Overall, the development of multimodal fusion has evolved from naive concatenation to transformer-driven, BEV-aligned, and memory-augmented strategies, yet cross-modal misalignment remains insufficiently addressed. This gap motivates our subspace-expert fusion framework, which explicitly disentangles modality-private, shared, and interaction-specific subspaces to achieve more robust and interpretable multi-modal HD map prediction.

III. METHODOLOGY

A. Problem Formulation and Method Overview

Problem Formulation. Given synchronized multi-view images $I = \{I_c\}_{c=1}^{N_{\text{cam}}}$ and a LiDAR point cloud L , our goal is to predict an HD map $\hat{\mathbf{Y}}$ in a BEV frame. Let modality encoders produce BEV features $F_I = E_I(I) \in \mathbb{R}^{H \times W \times C}$ and $F_L = E_L(L) \in \mathbb{R}^{H \times W \times C}$. Our method introduces a fusion module \mathcal{F}_θ (detailed in Sec. III-B) that operates purely in BEV and is followed by task heads Dec to decode vectorized map elements. The learning objective on dataset $\mathcal{D} = \{(I^{(n)}, L^{(n)}, \mathbf{Y}^{(n)})\}_{n=1}^N$ is

$$\min_{\theta, \phi} \frac{1}{N} \sum_{n=1}^N \mathcal{L}_{\text{task}} \left(Dec_\phi(\mathcal{F}_\theta(F_I^{(n)}, F_L^{(n)})), \mathbf{Y}^{(n)} \right), \quad (1)$$

where $\mathcal{L}_{\text{task}}$ is the standard set-based classification and geometric regression loss used for HD map elements.

Method Overview. The overall architecture of our method is shown in Fig. 1. To enable precise and adaptive cross-modal fusion, we decompose BEV features into four subspaces: LiDAR-private, Image-private, Shared, and Interaction spaces. A cell-wise, uncertainty-aware gating module then reweights these experts to form the final prediction. During training, we clarify expert roles and improve single-modality behavior by replacing one modality with a surrogate sampled from its empirical BEV feature distribution, creating realistic masked variants without disrupting feature statistics. Learning is further stabilized by a lightweight balance regularizer that prevents expert starvation. At inference, a single intact forward aggregates expert outputs via uncertainty-aware gating, without masking or auxiliary regularizers.

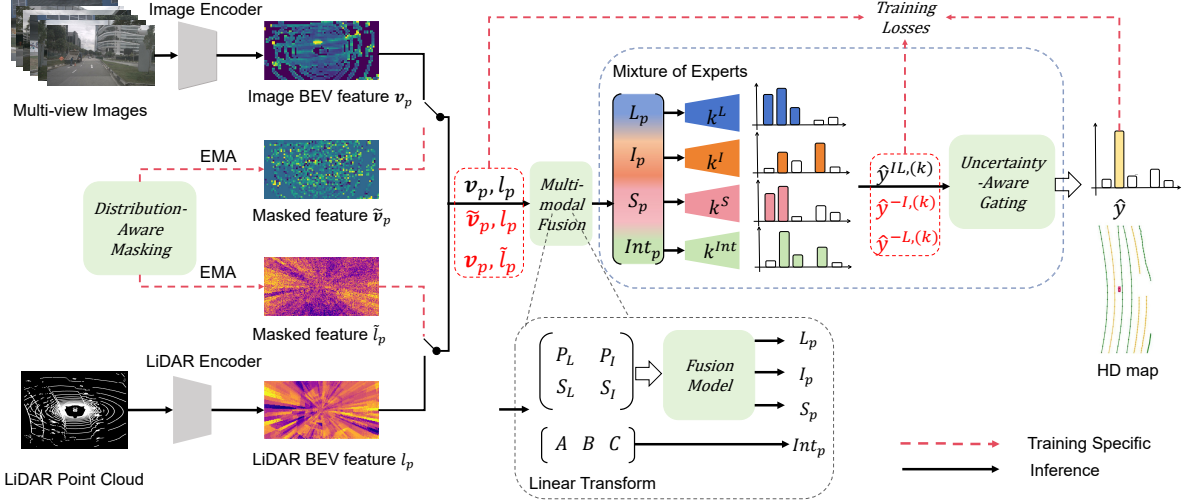


Fig. 1: **Overall architecture of SEF-MAP framework.** The system takes multi-view images and LiDAR point clouds as inputs, which are encoded into BEV features through respective encoders. These features are then decomposed into four semantic subspaces (LiDAR-private, Image-private, Shared, and Interaction) via linear transformations. During training, distribution-aware masking creates surrogate features using EMA statistics to simulate modality degradation scenarios, with specialization losses enforcing expert roles. During inference, only the intact forward pass is performed without masking. A mixture of experts processes each subspace, and an uncertainty-aware gating mechanism dynamically weights expert outputs based on their predicted variance to generate the final HD map prediction.

B. Subspace-Decomposed Multimodal Fusion

For each BEV cell $p \in \mathcal{P}$, we fuse the LiDAR BEV feature $\ell_p \in \mathbb{R}^C$ and the image BEV feature $\mathbf{v}_p \in \mathbb{R}^C$ by projecting them into four semantically distinct subspaces: *LiDAR* (\mathbf{L}_p), *Image* (\mathbf{I}_p), *Shared* (\mathbf{S}_p), and *Interaction* (\mathbf{Int}_p). With learnable linear maps $\mathbf{P}_L, \mathbf{P}_I, \mathbf{S}_L, \mathbf{S}_I \in \mathbb{R}^{C \times C}$ and a low-rank bilinear form parameterized by $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{R \times C}$ and $\mathbf{C} \in \mathbb{R}^{C \times R}$ with $R \ll C$, where C is the BEV channel width, R is the interaction rank. We first compute modality-private and shared projections:

$$\begin{aligned} \mathbf{u}_p^L &= \mathbf{P}_L \ell_p, & \mathbf{u}_p^I &= \mathbf{P}_I \mathbf{v}_p, \\ \mathbf{r}_p^L &= \mathbf{S}_L \ell_p, & \mathbf{r}_p^I &= \mathbf{S}_I \mathbf{v}_p, \end{aligned} \quad (2)$$

where \mathbf{u}_p^L preserves LiDAR-specific geometric cues such as range and 3D structure that remain stable under illumination changes or low-light conditions, while \mathbf{u}_p^I retains image-specific semantic cues such as appearance and road markings. The features \mathbf{r}_p^L and \mathbf{r}_p^I represent modality-shared evidence that is consistently observable across both LiDAR and camera, such as lane continuity.

To regulate subspace features, we introduce an auxiliary *space loss* defined as

$$\mathcal{L}_{\text{space}} = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} \left(\mathcal{L}_{\text{uni}} + \mathcal{L}_{\text{shr}} + \mathcal{L}_{\text{int}} \right), \quad (3)$$

where three components operate at each BEV cell p . The *private decorrelation loss* \mathcal{L}_{uni} enforces independence between private and opposite-modality features using HSIC:

$$\mathcal{L}_{\text{uni}} = \text{HSIC}(\mathbf{u}_p^L, \mathbf{v}_p) + \text{HSIC}(\mathbf{u}_p^I, \ell_p). \quad (4)$$

The *shared alignment loss* \mathcal{L}_{shr} encourages consistency between the shared projections \mathbf{r}_p^L and \mathbf{r}_p^I via negative correlation:

$$\mathcal{L}_{\text{shr}} = -\text{corr}(\mathbf{r}_p^L, \mathbf{r}_p^I). \quad (5)$$

The *interactive contrastive loss* \mathcal{L}_{int} strengthens the cross-modal sensitivity of the interaction subspace \mathbf{Int}_p through InfoNCE:

$$\mathcal{L}_{\text{int}} = -\log \frac{\exp \phi(\mathbf{Int}_p, \ell_p, \mathbf{v}_p)}{\sum_{p'} \exp \phi(\mathbf{Int}_p, \ell_{p'}, \mathbf{v}_{p'})}. \quad (6)$$

Together, these regularizers disentangle private, shared, and interaction subspaces, thereby improving modality robustness and cross-modal complementarity.

Based on these projections, four subspaces can be constructed as:

$$\begin{aligned} \mathbf{L}_p &= f_{\text{fusion}}^L([\mathbf{u}_p^L \parallel \mathbf{r}_p^L]), \\ \mathbf{I}_p &= f_{\text{fusion}}^I([\mathbf{r}_p^L \parallel \mathbf{u}_p^I]), \\ \mathbf{S}_p &= f_{\text{fusion}}^S([\mathbf{r}_p^L \parallel \mathbf{r}_p^I]), \\ \mathbf{Int}_p &= \mathbf{C}[(\mathbf{A} \ell_p) \odot (\mathbf{B} \mathbf{v}_p)]. \end{aligned} \quad (7)$$

where \odot denotes element-wise multiplication. In this formulation, \mathbf{L}_p and \mathbf{I}_p serve as the modality-inclined information, \mathbf{S}_p encodes common and modality-invariant information, and \mathbf{Int}_p captures complementary cross-modal interactions through multiplicative gating. The interaction space is particularly useful for resolving ambiguities caused by occlusions or weak signals in a single modality. For convenience, we denote $\mathbf{z}_p^{(k)} \in \{\mathbf{L}_p, \mathbf{I}_p, \mathbf{S}_p, \mathbf{Int}_p\}$ as the subspace features assigned to each expert $k \in K$, where $K = \{\mathbf{k}^L, \mathbf{k}^I, \mathbf{k}^S, \mathbf{k}^{Int}\}$. All linear maps are implemented as 1×1 BEV convolutions and operate independently at each BEV cell p .

C. Distribution-Aware Masking with Specialization Losses

Since either modality may be degraded by noise, occlusion, or sensor perturbations, we propose a *distribution-aware masked forward* strategy to encourage clear expert specialization and robustness to partial observations. To this end, we maintain per-channel means and variances of the BEV features, denoted as (μ_I, σ_I^2) for images and (μ_L, σ_L^2) for LiDAR, which are updated as exponential moving averages (EMA) during training to provide stable distribution estimates. When masking occurs, the suppressed modality is replaced with a surrogate feature sampled from its empirical distribution:

$$\tilde{\mathbf{v}}_p \sim \mathcal{N}(\mu_I, \text{diag}(\sigma_I^2)), \quad \tilde{\ell}_p \sim \mathcal{N}(\mu_L, \text{diag}(\sigma_L^2)). \quad (8)$$

During training, three types of forward passes are performed. The intact pass $\hat{\mathbf{y}}^{IL}$ uses both modalities (ℓ, \mathbf{v}) , the image-masked pass $\hat{\mathbf{y}}^{-I}$ uses LiDAR together with surrogate image features $(\ell, \tilde{\mathbf{v}})$, and the LiDAR-masked pass $\hat{\mathbf{y}}^{-L}$ uses camera features with surrogate LiDAR inputs $(\tilde{\ell}, \mathbf{v})$. These variants approximate unimodal conditions: private experts are encouraged to align with their respective modality, the shared expert is trained to remain consistent across both, and the interaction expert learns cues that weaken when either modality is missing. By sampling from empirical marginal distributions, the masked inputs preserve channel energy and normalization statistics, thereby avoiding out-of-distribution artifacts and providing informative gradients for optimization. At inference, only the intact forward pass is employed without masking.

To translate the behavior of masked passes into explicit learning signals, we design *specialization losses*. Specifically, we define per-cell losses that contrast intact and masked predictions, thereby clarifying the role of each expert.

We denote by $\hat{\mathbf{y}}^{IL,(k)}$, $\hat{\mathbf{y}}^{-I,(k)}$, and $\hat{\mathbf{y}}^{-L,(k)}$ the outputs of expert k under intact, image-masked, and LiDAR-masked conditions, respectively. Here, $d(\cdot, \cdot)$ is a task-consistent dissimilarity, $\gamma \in (0, 1]$ a balance coefficient, and $m > 0$ a margin.

For the LiDAR-private expert:

$$\mathcal{L}_p^L = d(\hat{\mathbf{y}}_p^{IL,(L)}, \hat{\mathbf{y}}_p^{-I,(L)}) - \gamma d(\hat{\mathbf{y}}_p^{IL,(L)}, \hat{\mathbf{y}}_p^{-L,(L)}), \quad (9)$$

while for the Image-private expert:

$$\mathcal{L}_p^I = d(\hat{\mathbf{y}}_p^{IL,(I)}, \hat{\mathbf{y}}_p^{-L,(I)}) - \gamma d(\hat{\mathbf{y}}_p^{IL,(I)}, \hat{\mathbf{y}}_p^{-I,(I)}). \quad (10)$$

The Shared expert is optimized to remain consistent under both maskings:

$$\mathcal{L}_p^S = d(\hat{\mathbf{y}}_p^{IL,(S)}, \hat{\mathbf{y}}_p^{-I,(S)}) + d(\hat{\mathbf{y}}_p^{IL,(S)}, \hat{\mathbf{y}}_p^{-L,(S)}), \quad (11)$$

and the Interaction expert emphasizes cross-modal complementarity:

$$\begin{aligned} \mathcal{L}_p^{Int} = & [m - d(\hat{\mathbf{y}}_p^{IL,(Int)}, \hat{\mathbf{y}}_p^{-I,(Int)})]_+ \\ & + [m - d(\hat{\mathbf{y}}_p^{IL,(Int)}, \hat{\mathbf{y}}_p^{-L,(Int)})]_+. \end{aligned} \quad (12)$$

Averaging across cells yields

$$\mathcal{L}_{\text{spec}} = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} (\mathcal{L}_p^L + \mathcal{L}_p^I + \mathcal{L}_p^S + \mathcal{L}_p^{Int}). \quad (13)$$

D. Uncertainty-Aware Gating and Stabilization

Subspace reliability varies across BEV cells, so the model should prefer the expert that is both informative and confident at each location. For cell $p \in \mathcal{P}$ and expert $k \in K$, the head H_k maps the subspace feature $\mathbf{z}_p^{(k)}$ to a mean prediction $\mu_p^{(k)} \in \mathbb{R}^D$ and a variance $\sigma_p^{2(k)} \in \mathbb{R}^D$. Here the variance is parameterized as $\log \sigma^2$ for stability. A lightweight gate g produces logits $\alpha_p^{(k)} \in \mathbb{R}$. We down-weight uncertain experts by subtracting their average predicted variance from the logits:

$$w_p^{(k)} = \text{softmax}_k(\alpha_p^{(k)} - \beta \bar{\sigma}_p^{2(k)}), \quad \bar{\sigma}_p^{2(k)} = \frac{1}{D} \sum_{d=1}^D \sigma_{p,d}^{2(k)}, \quad (14)$$

where $\beta > 0$ controls the penalty strength and D is the head output dimension. The final per-cell prediction is

$$\hat{\mathbf{y}}_p = \sum_{k \in K} w_p^{(k)} \mu_p^{(k)} \in \mathbb{R}^D. \quad (15)$$

Moreover, gating can collapse onto a single expert if left unconstrained, which is called expert starvation and harms specialization. Therefore, we regularize the average usage of each expert toward a uniform prior $1/|K|$,

$$\Omega_{\text{bal}} = \sum_{k \in K} \left(\bar{w}_k - \frac{1}{|K|} \right)^2, \quad \bar{w}_k = \frac{1}{|\mathcal{P}|} \sum_{p \in \mathcal{P}} w_p^{(k)}. \quad (16)$$

Together, Eqs. (14)–(16) yield gates that are cell-wise adaptive, variance-aware, and resistant to expert collapse.

Algorithm 1 Training and Inference for Subspace-Decomposed Fusion

Inputs : Batch $\{(I, L, \mathbf{Y})\}$; encoders E_I, E_L ; lifts T_I, T_L ; maps $\mathbf{P}_*, \mathbf{S}_*, \mathbf{A}, \mathbf{B}, \mathbf{C}$; heads $\{H_k\}_{k \in K}$; gate g ; EMA stats $(\mu_I, \sigma_I^2), (\mu_L, \sigma_L^2)$

Output : $\hat{\mathbf{y}}$

```

1 procedure TrainBatch
2    $F_I \leftarrow T_I(E_I(I)); F_L \leftarrow T_L(E_L(L))$  for  $p \in \mathcal{P}$  do
3      $\mathbf{u}_p^L, \mathbf{u}_p^I, \mathbf{u}_p^S, \mathbf{u}_p^{Int} \leftarrow$  Eq. (7); set  $\mathbf{z}_p^{(k)} (\mu_p^{(k)}, \sigma_p^{2(k)}) \leftarrow H_k(\mathbf{z}_p^{(k)})$ 
4     for  $k \in K$ 
5      $\{\alpha_p^{(k)}\} \leftarrow g(\{\mathbf{z}_p^{(k)}\}_k)$   $w_p^{(k)} \leftarrow$  Eq. (14)  $\hat{\mathbf{y}}^{IL} \leftarrow$  Eq. (15) Sample
6      $\tilde{\mathbf{v}}_p, \tilde{\ell}_p$  via Eq. (8)  $\triangleright$  compute  $\hat{\mathbf{y}}^{-I}, \hat{\mathbf{y}}^{-L}$ 
7     Compute  $\mathcal{L}$  using Eqs. (3)–(6), (9)–(13), and (16)–(18)
8      $\triangleright$  update model parameters and EMA stats
9
10 procedure Inference $_{I,L}$ 
11    $F_I \leftarrow T_I(E_I(I)); F_L \leftarrow T_L(E_L(L))$  for  $p \in \mathcal{P}$  do
12      $\mathbf{z}_p^{(k)} \leftarrow$  Eq. (7);  $\mu_p^{(k)} \leftarrow H_k(\mathbf{z}_p^{(k)})$  for  $k \in K$ 
13      $w_p^{(k)} \leftarrow$  Eq. (14); return  $\hat{\mathbf{y}} = \sum_{k \in K} w_p^{(k)} \mu_p^{(k)} \triangleright$  Eq. (15)
```

E. Training and Inference Pipeline

Loss Definition. Aggregating per-cell predictions yields the set-level output $\hat{\mathbf{Y}}^{IL}$. The task loss compares intact predictions against ground truth:

$$\mathcal{L}_{\text{task}} = \mathcal{L}_{\text{task}}(\hat{\mathbf{Y}}^{IL}, \mathbf{Y}). \quad (17)$$

Finally, the complete training objective integrates all components:

$$\mathcal{L} = \mathcal{L}_{\text{task}} + \lambda_{\text{spec}} \mathcal{L}_{\text{spec}} + \lambda_{\text{space}} \mathcal{L}_{\text{space}} + \lambda_{\text{bal}} \Omega_{\text{bal}}, \quad (18)$$

TABLE I: **Comparisons with state-of-the-art methods on nuScenes val set.** Our method outperforms all existing approaches in both single-class APs and overall mAP.

Method	Modality	Backbone	Epochs	AP_{ped}	$AP_{divider}$	$AP_{boundary}$	mAP
HMapNet [41]	C	E-B0	30	14.4	21.7	33.0	23.0
HMapNet [41]	L	PointPillars	30	10.4	24.1	37.9	24.1
HMapNet [41]	C&L	E-B0 & PointPillars	30	16.3	29.6	46.7	31.0
VectorMapNet [42]	C	R50	110	36.1	47.3	39.3	40.9
VectorMapNet [42]	L	PointPillars	110	25.7	37.6	38.6	34.0
VectorMapNet [42]	C&L	R50 & PointPillars	110	37.6	50.5	47.5	45.2
MapTR [43]	C	R50	24	46.3	51.5	53.1	50.3
MapTR [43]	L	SECOND	24	48.5	53.7	64.7	55.6
MapTR [43]	C&L	R50 & SECOND	24	55.9	62.3	69.3	62.5
SEF-Map (Ours)	C&L	R50 & SECOND	24	61.6	66.7	71.8	66.7(+4.2)

TABLE II: **Results on Argoverse2 dataset.** † denotes our re-implementation following the setting in the paper.

Method	Modality	Backbone	Epochs	AP_{ped}	$AP_{divider}$	$AP_{boundary}$	mAP
HMapNet [41]	C	E-B0	30	13.1	5.7	37.6	18.8
VectorMapNet [42]	C	R50	110	38.3	36.1	39.2	37.9
MapTR [†] [43]	C	R50	6	58.7	59.3	60.3	59.4
MapTR [†] [43]	C&L	R50 & SECOND	6	65.1	61.6	75.1	67.3
SEF-Map (Ours)	C&L	R50 & SECOND	6	70.7	66.7	78.8	72.1(+4.8)

where $\lambda_{spec}, \lambda_{space}, \lambda_{bal} > 0$ balance the contributions of different terms.

Training. Each mini-batch mixes intact and masked conditions so the gate learns when to trust private, shared, or interaction evidence. Inputs are encoded and lifted to BEV. For each cell we compute the four subspace features using Eq. (7), and the expert heads return mean and variance estimates. The gate produces logits and we apply the variance penalty in Eq. (14) to obtain weights; the intact prediction then follows from Eq. (15). Surrogate features are sampled with Eq. (8) to create two masked passes, and the same forward path yields their predictions. We evaluate the objective in Eq. (18) and finally update network parameters by backpropagation and refresh the exponential moving averages used for distribution-aware masking.

Inference. At test time, we perform a single intact forward pass without any masking or auxiliary objectives. For each BEV cell, the model first computes the corresponding subspace features, which are then processed by the expert heads to produce both mean predictions and associated variances. The gating module subsequently generates weights according to Eq. (14), allowing reliable experts to contribute more strongly. The final output is obtained as a weighted mixture of expert predictions, as defined in Eq. (15). Since no additional masking or auxiliary losses are involved during inference, the runtime remains comparable to that of a standard BEV fusion network.

IV. EXPERIMENTS

A. Experimental Setup

Dataset. We conducted experiments on the nuScenes and Argoverse2 datasets. The nuScenes dataset contains 1000 driving sequences, each with six camera streams and LiDAR sweeps annotated at 2Hz. We follow the preprocessing pipeline in MapTR [43] to ensure consistent data preparation. In our experiments, we focus on three representative map elements: pedestrian crossings, lane dividers, and road boundaries. The Argoverse2 dataset also provides 1000 sequences, partitioned into 700 for training, 150 for validation, and 150 for testing. Each sequence spans 15 seconds, with synchronized 20Hz camera images and 10Hz LiDAR sweeps. For both datasets, we apply spatial alignment, temporal synchronization, and feature normalization to enable reliable multi-modal fusion during training. Owing to their diverse urban and suburban scenes, these datasets serve as strong benchmarks for evaluating the generalization ability of our method.

Evaluation Metrics. We evaluate the quality of the predicted HD maps using average precision (AP). To assess the alignment between predicted and ground-truth map elements, we compute the Chamfer distance ($D_{Chamfer}$) at multiple thresholds ($\tau \in T$, where $T = \{0.5, 1.0, 1.5\}$ meters). The final AP score is averaged over all these thresholds, providing a comprehensive measure of the model’s performance across different levels of spatial accuracy.

Implementation Details. We train SEF-MAP on 8 NVIDIA RTX 4090 GPUs using the AdamW optimizer. The learning rate is set to $6e-4$ with step decay, and the mini-batch size is 16. The backbone uses ResNet50 for image features and SECOND for LiDAR features. The overall loss function consists of task loss, space loss, specialization loss, and balance regularization, with weights $\lambda_{\text{task}} = 1.0$, $\lambda_{\text{space}} = 0.1$, $\lambda_{\text{spec}} = 0.5$, and $\lambda_{\text{bal}} = 0.05$. To ensure consistency, the sampled forward passes adopt exactly the same input configuration as the intact inputs, so that the gating mechanism learns under comparable conditions.

B. Comparison with SOTA Methods

We compare **SEF-MAP** with several state-of-the-art models, including HDMaPNet [41], VectorMapNet [42], and MapTR [43], on the nuScenes and Argoverse2 datasets, as shown in Tab. I and Tab. II. **SEF-MAP** consistently outperforms all baselines across both datasets and all map element categories: on nuScenes, the most notable improvement is a +5.7% gain in AP_{ped} compared to MapTR, along with substantial improvements of +4.4% in AP_{divider} and +2.5% in AP_{boundary} , resulting in an overall mAP improvement of +4.2%. On Argoverse2, **SEF-MAP** achieves even more significant gains, with the largest boost in overall mAP of +4.8%, demonstrating strong generalization capability across different datasets and driving scenarios. Notably, the consistent improvements across all three map element types (pedestrian crossings, lane dividers, and road boundaries) indicate that our framework effectively captures both fine-grained semantic details and geometric structures. These results highlight the effectiveness of our subspace decomposition and uncertainty-aware gating modules in handling diverse environmental conditions and sensor configurations. The substantial performance gains further indicate that the superiority of multi-modal designs over single-modality counterparts is systematic rather than incidental, underscoring the importance of fully exploiting modality-specific representations while maintaining robust cross-modal interactions to guide the learning process for reliable HD map construction.

C. Ablation Studies

Analysis on Different Modules. To systematically evaluate the effectiveness of each module in our proposed model, we train the model by removing each component one at a time, as presented in Tab. III. We design the following configurations: We design the following ablation models: (1) **w/o DAM + w/o UAG**: we remove DAM and replace UAG with an MLP; (2) **w/o SD + w/o UAG**: we remove SD and replace UAG with an MLP; (3) **w/o SD + w/o DAM**: we remove both SD and DAM modules; (4) **w/o UAG**: we replace *uncertainty-aware gating (UAG)* with an MLP for computing expert weights; (5) **w/o DAM**: we remove the *distribution-aware masking (DAM)* module; (6) **w/o SD**: we remove the *subspace decomposition (SD)* module from the model;

The results show that introducing SD and DAM brings the most substantial gains, with mAP improving by about

TABLE III: **Ablation on subspace decomposition (SD), distribution-aware masking (DAM), and uncertainty-aware gating (UAG).** The full model combining all three components achieves the state-of-the-art performance, outperforming individual modules and pairwise combinations on nuScenes val (24 epochs).

SD	DAM	UAG	Mod.	AP_{ped}	AP_{divider}	AP_{boundary}	mAP
×	×	×	C&L	55.9	62.3	69.3	62.5
✓	×	×	C&L	56.8	63.1	70.2	63.4
×	✓	×	C&L	57.7	64.6	71.0	64.4
×	×	✓	C&L	56.1	62.8	69.7	62.9
✓	✓	×	C&L	59.1	66.7	72.0	65.9
✓	×	✓	C&L	59.2	64.9	70.7	64.9
×	✓	✓	C&L	57.8	65.3	72.2	65.1
✓	✓	✓	C&L	61.6	66.7	71.8	66.7

TABLE IV: **Expert group ablation configurations and results.** The full model achieves the best performance, significantly outperforming configurations with only private experts (L+I) or only cross-modal experts (S+Int).

Expert Setting	AP_{ped}	AP_{divider}	AP_{boundary}	mAP
MapTR [43] (baseline)	55.9	62.3	69.3	62.5
Only L+I (Private)	57.7	64.2	69.4	63.8
Only S+Int (Cross-modal)	54.6	60.5	66.7	60.6
SEF-Map (Full)	61.6	66.7	71.8	66.7

3.4% compared to the baseline. UAG alone yields only a marginal +0.4% gain, but when combined with SD and DAM it provides an additional +0.8% improvement, showing that UAG is most effective when enhancing their complementary features. These findings confirm the complementary roles of the three modules, with SD and DAM providing the strongest improvements and UAG refining the fusion process.

Analysis on Expert Group Ablation. Additionally, we perform an expert group ablation to evaluate the importance of different expert configurations, as shown in Tab. IV. We design the following configurations: (1) **Private Experts Only (L+I)**: using only the *private experts from LiDAR and Image*; (2) **Cross-Modal Experts Only (S+Int)**: using only the *cross-modal experts*. The results show that using only private experts (L+I) achieves slightly better performance than the baseline MapTR (+1.3% mAP), with improvements particularly evident in AP_{ped} (+1.8%) and AP_{divider} (+1.9%), suggesting that modality-specific features are valuable for capturing fine-grained semantic and geometric details. Using only cross-modal experts (S+Int) leads to even weaker performance (-1.9% compared to baseline), with notable drops in AP_{boundary} (-2.6%) and AP_{ped} (-1.2%), confirming that cross-modal experts alone cannot capture all the necessary modality-specific cues for accurate HD map construction, such as LiDAR’s precise depth information or camera’s rich texture details. The **SEF-MAP**, combining SD, DAM, and UAG, achieves the highest performance

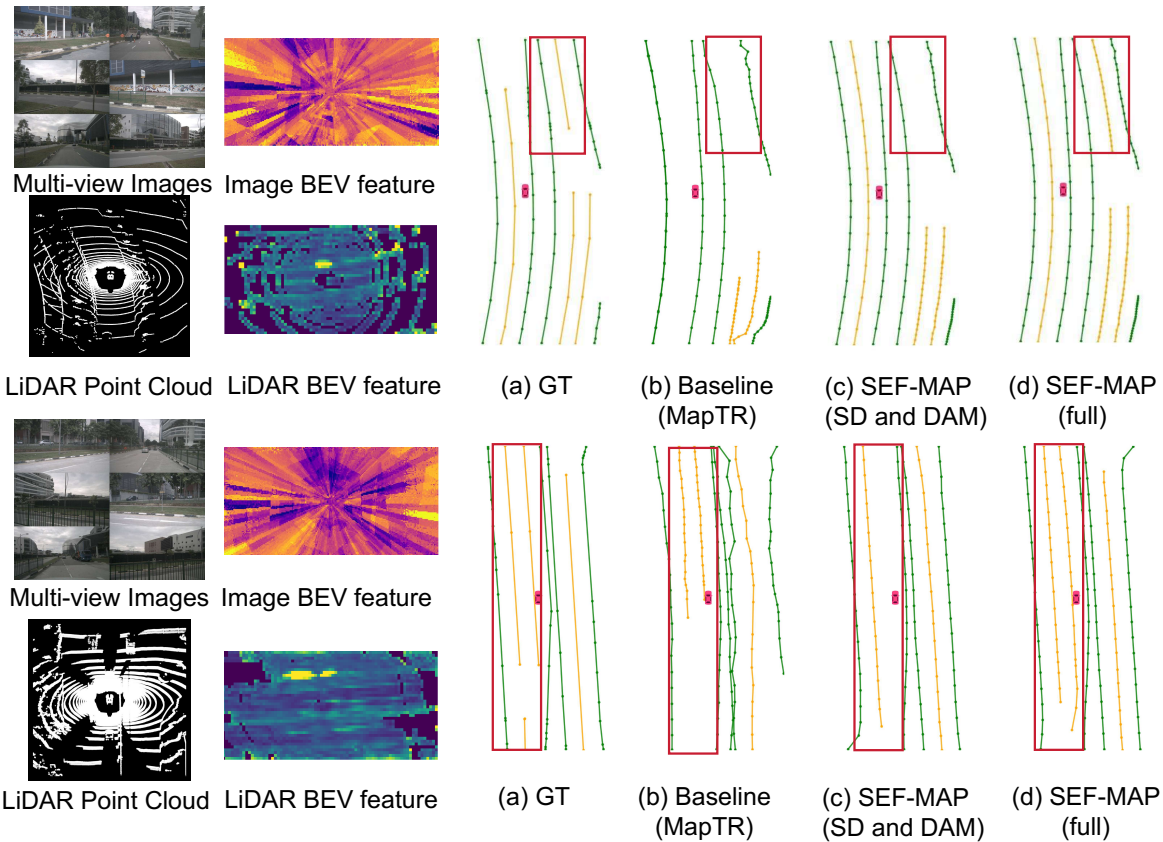


Fig. 2: **Qualitative results on nuScenes.** We present two sample scenes from nuScenes: (a) Ground Truth. (b) Baseline (MapTR). (c) SEF-MAP (SD and DAM). (d) **SEF-MAP (full)**. The red boxes highlight challenging regions where our method demonstrates significant improvements over the baseline. In these highlighted areas, the baseline MapTR produces noticeable errors and incomplete map predictions. The full SEF-MAP model with uncertainty-aware gating (d) achieves the most precise vectorized map reconstruction, particularly in the challenging regions marked by red boxes, demonstrating the effectiveness of our subspace decomposition and adaptive fusion strategy.

(+4.2% mAP), with consistent improvements across all map elements, demonstrating that the synergy between private and cross-modal experts is crucial for leveraging both modality-specific strengths and cross-modal interactions to achieve robust multimodal fusion.

D. Qualitative Results

In Fig 2, we show qualitative results on two sample scenes from the nuScenes dataset with both LiDAR and camera inputs. We compare the predicted vectorized HD map results of different models, including the baseline MapTR [43], our model with SD and DAM, and the full model (with SD, DAM, and UAG). The baseline model produces significant errors in map prediction, particularly evident in the red-boxed regions where lane boundaries are misaligned or completely missing. Our progressive model variants show increasingly better performance, with the SD and DAM combination providing substantial improvements in lane detection and road topology preservation. The full SEF-MAP model achieves the best results, demonstrating superior alignment with ground truth and successful reconstruction of complex lane structures and road boundaries that were problematic for the baseline. These visual results confirm that the synergistic

combination of SD, DAM, and UAG effectively enhances multi-modal fusion, ultimately producing more reliable and accurate HD map predictions.

V. CONCLUSION

In this work, we propose **SEF-MAP**, a Subspace-Expert Fusion framework for robust multi-modal HD map prediction. **SEF-MAP** decomposes BEV features into LiDAR-private, Image-private, Shared, and Interaction subspaces, each handled by dedicated experts to balance modality-specific information and cross-modal consensus. An uncertainty-aware gating mechanism adaptively regulates expert contributions, while distribution-aware masking with specialization constraints improves robustness under degraded conditions. Extensive experiments on nuScenes and Argoverse2 datasets demonstrate the effectiveness of our approach, achieving significant improvements of +4.2% and +4.8% mAP respectively over state-of-the-art baselines. The principled design of expert specialization and adaptive fusion not only enhances performance but also provides interpretable insights into modality contributions. Overall, **SEF-MAP** offers an effective solution for reliable multi-modal HD map prediction under diverse and challenging real-world scenarios.

REFERENCES

- [1] X. Hao, Y. Zhao, Y. Ji *et al.*, “What really matters for robust multi-sensor hd map construction?” *arXiv preprint arXiv:2507.01484*, 2025.
- [2] G. Elghazaly, R. Frank, S. Harvey *et al.*, “High-definition maps: Comprehensive survey, challenges, and future perspectives,” *IEEE Open Journal of Intelligent Transportation Systems*, vol. 4, pp. 527–550, 2023.
- [3] Q. Li, Y. Wang, Y. Wang *et al.*, “Hdmapnet: An online hd map construction and evaluation framework,” in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 4628–4634.
- [4] B. Wijaya, K. Jiang, M. Yang *et al.*, “High definition map mapping and update: A general overview and future directions,” *arXiv preprint arXiv:2409.09726*, 2024.
- [5] X. Hao, M. Wei, Y. Yang, H. Zhao, H. Zhang, Y. Zhou, Q. Wang, W. Li, L. Kong, and J. Zhang, “Is your hd map constructor reliable under sensor corruptions?” *Advances in Neural Information Processing Systems*, pp. 22 441–22 482, 2024.
- [6] X. Hao, Y. Diao, M. Wei, Y. Yang, P. Hao, R. Yin, H. Zhang, W. Li, S. Zhao, and Y. Liu, “Mapfusion: A novel bev feature fusion network for multi-modal map construction,” *Information Fusion*, p. 103018, 2025.
- [7] X. Hao, G. Liu, Y. Zhao, Y. Ji, M. Wei, H. Zhao, L. Kong, R. Yin, and Y. Liu, “Msc-bench: Benchmarking and analyzing multi-sensor corruption for driving perception,” *arXiv preprint arXiv:2501.01037*, 2025.
- [8] Z. Bao, S. Hossain, H. Lang *et al.*, “High-definition map generation technologies for autonomous driving,” *arXiv preprint arXiv:2206.05400*, 2022.
- [9] X. Tang, K. Jiang, M. Yang *et al.*, “High-definition maps construction based on visual sensor: A comprehensive survey,” *IEEE Transactions on Intelligent Vehicles*, 2023.
- [10] H. Dong, W. Gu, X. Zhang *et al.*, “Superfusion: Multilevel lidar-camera fusion for long-range hd map generation,” in *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2024, pp. 9056–9062.
- [11] M. E. Kalfaoglu, H. I. Ozturk, O. Kilinc, and A. Temizel, “Topomaskv2: Enhanced instance-mask-based formulation for the road topology problem,” in *CVPR*, 2024, pp. 17 853–17 862.
- [12] C. Zhang, Q. Song, F. Li *et al.*, “Hybrimap: Hybrid clues utilization for effective vectorized hd map construction,” *arXiv preprint arXiv:2404.11155*, 2024.
- [13] Y. Zhou, H. Zhang, J. Yu *et al.*, “Himap: Hybrid representation learning for end-to-end vectorized hd map construction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 15 396–15 406.
- [14] L. Qiao, Y. Zheng, P. Zhang *et al.*, “Machmap: End-to-end vectorized solution for compact hd-map construction,” *arXiv preprint arXiv:2306.10301*, 2023.
- [15] L. Qiao, W. Ding, X. Qiu, and C. Zhang, “End-to-end vectorized hd-map construction with piecewise bezier curve,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 13 218–13 228.
- [16] L. Peng, Z. Chen, Z. Fu, P. Liang, and E. Cheng, “Bevsegformer: Bird’s eye view semantic segmentation from arbitrary camera rigs,” in *IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023, pp. 5924–5932.
- [17] Z. Qin, J. Chen, C. Chen *et al.*, “Uniformer: Unified multi-view fusion transformer for spatial-temporal representation in birds-eye-view,” *arXiv preprint arXiv:2207.08536*, 2022.
- [18] Y. Zhang, Z. Zhu, W. Zheng *et al.*, “Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving,” *arXiv preprint arXiv:2205.09743*, 2022.
- [19] Z. Li, W. Wang, H. Li *et al.*, “Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers,” in *European Conference on Computer Vision*, 2022, pp. 1–18.
- [20] C. Yang, Y. Chen, H. Tian *et al.*, “Bevformer v2: Adapting modern image backbones to bird’s-eye-view recognition via perspective supervision,” *arXiv preprint arXiv:2211.10439*, 2022.
- [21] R. Zhou, C. Dai, W. Jiang *et al.*, “Supermapnet for long-range and high-accuracy vectorized hd map construction,” *arXiv preprint arXiv:2505.13856*, 2025.
- [22] X. Hao, R. Li, H. Zhang *et al.*, “Mapdistill: Boosting efficient camera-based hd map construction via camera-lidar fusion model distillation,” in *European Conference on Computer Vision*, 2024, pp. 166–183.
- [23] J. Song, X. Chen, L. Lu *et al.*, “Memfusionmap: Working memory fusion for online vectorized hd map construction,” in *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025, pp. 9248–9257.
- [24] Y. Liu, T. Yuan, Y. Wang *et al.*, “Vectormapnet: End-to-end vectorized hd map learning,” in *International Conference on Machine Learning*, 2023, pp. 22 352–22 369.
- [25] Z. Zhou, X. Feng, S. Di *et al.*, “A lidar mapping system for robot navigation in dynamic environments,” *IEEE Transactions on Intelligent Vehicles*, 2023.
- [26] W. Ali, P. Jensfelt, and T. M. Nguyen, “Hd-maps as prior information for globally consistent mapping in gps-denied environments,” in *IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*, 2024, pp. 2145–2151.
- [27] X. Hao, L. Kong, R. Yin, P. Wang, J. Zhang, Y. Diao, and S. Zhao, “Safemap: Robust hd map construction from incomplete observations,” in *Forty-second International Conference on Machine Learning*.
- [28] X. Hao, H. Zhang, Y. Yang, and *et al.*, “Mbfusion: A new multi-modal bev feature fusion method for hd map construction,” in *IEEE International Conference on Robotics and Automation*, 2024, pp. 15 922–15 928.
- [29] J. Chen and A. Zhang, “Hgmf: heterogeneous graph-based fusion for multimodal data with incompleteness,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 1295–1305.
- [30] Q. Zhang, Y. Wei, Z. Han *et al.*, “Multimodal fusion on low-quality data: A comprehensive survey,” *arXiv preprint arXiv:2404.18947*, 2024.
- [31] X. Chen, H. Ma, J. Wan *et al.*, “Multi-view 3d object detection network for autonomous driving,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1907–1915.
- [32] K. Zhao, L. Liu, Y. Meng *et al.*, “Feature deep continuous aggregation for 3d vehicle detection,” *Applied Sciences*, vol. 9, no. 24, p. 5397, 2019.
- [33] J. Xiao, S. Wang, J. Zhou *et al.*, “Mim: High-definition maps incorporated multi-view 3d object detection,” *IEEE Transactions on Intelligent Transportation Systems*, 2024.
- [34] X. Han, S. Chen, Z. Fu *et al.*, “Multimodal fusion and vision-language models: A survey for robot vision,” *arXiv preprint arXiv:2504.02477*, 2025.
- [35] X. Bai, Z. Hu, X. Zhu *et al.*, “Transfusion: Robust lidar-camera fusion for 3d object detection with transformers,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 1090–1099.
- [36] Z. Liu, H. Tang, A. Amini *et al.*, “Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation,” in *IEEE International Conference on Robotics and Automation*, 2023, pp. 2774–2781.
- [37] Q. Jiang and H. Sun, “Semanticbevfusion: Rethinking lidar-camera fusion in unified bird’s-eye view representation for 3d object detection,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 5707–5714.
- [38] T. Liang, H. Xie, K. Yu *et al.*, “Bevfusion: A simple and robust lidar-camera fusion framework,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 10 421–10 434, 2022.
- [39] Y. Xie, C. Xu, M. J. Rakotosaona *et al.*, “Sparsefusion: Fusing multi-modal sparse representations for multi-sensor 3d object detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 17 591–17 602.
- [40] K. Wu, C. Yang, and Z. Li, “Interactionmap: Improving online vectorized hdmap construction with interaction,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025, pp. 17 176–17 186.
- [41] Q. Li, Y. Wang, Y. Wang, and H. Zhao, “Hdmapnet: An online hd map construction and evaluation framework,” in *IEEE ICRA*, 2022, pp. 4628–4634.
- [42] Y. Liu, T. Yuan, Y. Wang, Y. Wang, and H. Zhao, “Vectormapnet: End-to-end vectorized hd map learning,” in *ICML*, 2023, pp. 22 352–22 369.
- [43] B. Liao, S. Chen, X. Wang, T. Cheng, Q. Zhang, W. Liu, and C. Huang, “Maptr: Structured modeling and learning for online vectorized hd map construction,” in *ICLR*, 2023.