

Background *Fades*, Foreground *Leads*: Curriculum-Guided Background Pruning for Efficient Foreground-Centric Collaborative Perception

Yuheng Wu¹, Xiangbo Gao², Quang Tau¹, Zhengzhong Tu² and Dongman Lee¹

Abstract—Collaborative perception enhances the reliability and spatial coverage of autonomous vehicles by sharing complementary information across vehicles, offering a promising solution to long-tail scenarios that challenge single-vehicle perception. However, the bandwidth constraints of vehicular networks make transmitting the entire feature map impractical. Recent methods, therefore, adopt a foreground-centric paradigm, transmitting only predicted foreground-region features while discarding the background, which encodes essential context. We propose *FadeLead*, a foreground-centric framework that overcomes this limitation by learning to encapsulate background context into compact foreground features during training. At the core of our design is a curricular learning strategy that leverages background cues early on but progressively prunes them away, forcing the model to internalize context into foreground representations without transmitting background itself. Extensive experiments on both simulated and real-world benchmarks show that *FadeLead* outperforms prior methods under different bandwidth settings, underscoring the effectiveness of context-enriched foreground sharing. All artifacts are available at <https://wyhallenwu.github.io/FadeLead/>.

I. INTRODUCTION

Collaborative perception has emerged as a cornerstone in advancing the safety and reliability of autonomous driving systems. By enabling multiple vehicles to share complementary sensory information, it overcomes critical limitations of single vehicle perception, such as occlusions, restricted fields of view, and blind spots [1–11]. Through multi-agent cooperation, collaborative perception systems achieve broader spatial coverage and improved perception robustness.

Despite these benefits, real-world deployment must operate under stringent communication constraints imposed by vehicular networks. A central challenge, therefore, is determining *what subset of features* to transmit to maximize perception utility while minimizing bandwidth consumption.

To address this challenge, recent approaches [4, 6, 12] have predominantly adopted a *foreground-centric paradigm*. The intuition is straightforward: object-centric regions (*e.g.*, vehicles, pedestrians) are compact in nature, highly task-relevant, and therefore prioritize them for transmission. Typical pipelines learn a confidence map to estimate the foreground likelihood of each BEV cell, and only the top- k high-confidence regions are transmitted. However, this foreground-centric paradigm implicitly treats background as redundant and assumes that confident foreground regions alone are sufficient for robust collaborative perception. This raises a fundamental and underexplored question:

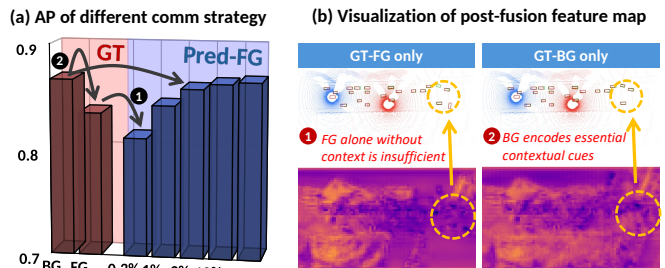


Fig. 1: (a) Average Precision (AP) of Where2Comm [4] on OPV2V [13] under different communication strategies and (b) visualization of the post-fusion feature map.

Is background, which often constitutes the majority of the scene yet is typically discarded, truly redundant?

Motivation. To ground our study, we first revisit the representative method Where2Comm [4] and disentangle the roles of three feature types: (1) Predicted Foreground (Pred-FG): regions selected by confidence scores. (2) Ground-Truth Foreground (GT-FG): precise object regions from ground-truth annotations. (3) Ground-Truth Background (GT-BG): all remaining regions complementing GT-FG. We design two experiments to investigate their contribution:

Experiment 1 → **Insight 1**: **GT-FG Only.** We replace Pred-FG with GT-FG, transmitting only precise object regions. Although these regions occupy merely $\sim 0.3\%$ (on average across the test split) of the BEV plane, they outperform the same fraction of Pred-FG, indicating that accurate object localization is more reliable. However, as the selection ratio increases (*e.g.*, $0.3\% \rightarrow 90\%$ blue bars in Fig. 1a), the performance of Pred-FG continues to improve. These results reveal that **foreground alone, even when perfectly localized, remains insufficient**, as it neglects inter-object dependencies and broader scene semantics.

Experiment 2 → **Insight 2**: **GT-BG Only.** Conversely, we mask out all GT-FG regions and transmit only the background region features. Surprisingly, as shown in Fig. 1, the GT-BG strategy not only surpasses the GT-FG counterpart but also achieves performance comparable to sharing nearly the entire BEV feature map. Moreover, as shown in Fig. 1b, they yield cleaner post-fusion representations. These results demonstrate that **background is not redundant but carries rich contextual cues** that are critical for disambiguation, robustness under occlusions, and holistic scene understanding.

Together, these insights motivate a new perspective on efficient, bandwidth-constrained collaborative perception:

Foreground should not stand alone. It must be enriched by the essential context encoded in the background.

¹KAIST {yuhengwu, quangntau1223, dlee}@kaist.ac.kr

²Texas A&M University {xiangbog, tzz}@tamu.edu

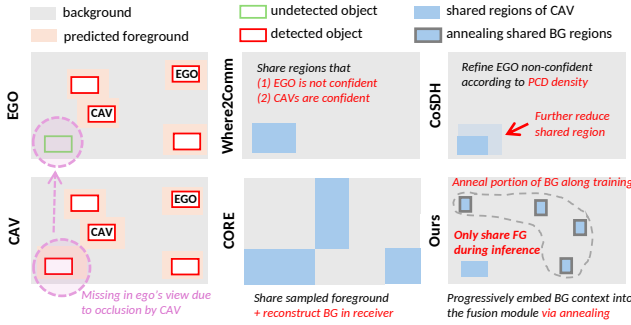


Fig. 2: Comparison of different foreground-centric sharing strategy.

Proposed Framework. We propose FadeLead, a collaborative perception framework that *remains foreground-centric* but addresses the incompleteness of foreground-only sharing. The core idea is to enrich compact foreground representations with critical background context, thereby preserving communication efficiency while enhancing perception robustness. To this end, FadeLead sharpens shared foreground features through three complementary modules: (1) *Foreground–Context Attention (FCA)* enriches foreground features by querying the full scene for complementary cues. (2) *Curricular Background Pruning (CBP)*, our major contribution, selectively mines informative background during training and progressively prunes it with an annealing schedule, forcing the model to internalize background-derived context into compact foreground features. At inference, only the enriched foreground is transmitted, enabling efficient yet context-aware collaboration. (3) *Foreground Amplification Fusion (FAF)* selectively enriches and amplifies salient ego foreground features by fusing complementary cues from local observations and received neighbor features, ensuring robust perception under strict bandwidth constraints.

Contributions. We summarize our contributions as follows:

- We conduct a systematic investigation on the role of background in collaborative perception. Our motivating study reveals that background region features, typically discarded by foreground-centric strategies, encode essential contextual cues for robust perception.
- We propose FadeLead, a curricular training framework that *remains foreground-centric* but strengthens it by encapsulating informative background context into compact shared foreground features, thereby improving both bandwidth efficiency and contextual completeness.
- We demonstrate through extensive experiments on simulated and real-world scenarios that FadeLead achieves superior detection accuracy under the same bandwidth budget, and remains highly effective even with extremely low selection ratios (*i.e.*, 1%).

II. RELATED WORK

Collaborative Perception. The information-sharing mechanisms in collaborative perception can be broadly categorized into three paradigms: (1) Early fusion [14, 15], which transmits raw sensor data (*e.g.*, images or point clouds) between vehicles. While this approach offers high perception accuracy, it incurs prohibitive communication costs due to

the large data volumes of raw inputs. (2) Late fusion [12, 13], which exchanges only the final detection results from each vehicle. This approach is bandwidth-efficient but suffers from higher latency (since full-model inference must be completed before sharing) and low update frequency of detection outputs. (3) Intermediate fusion [4, 5, 12, 16–18], which has emerged as the dominant strategy by striking a balance between these extremes. In this paradigm, intermediate feature representations, often the BEV feature, are transmitted, allowing for richer information sharing at a moderate communication cost. A common assumption, initiated by Where2Comm [4], is that foreground region features (*e.g.*, vehicles, pedestrians) provide the most informative cues while keeping transmission costs manageable.

Foreground-Centric Sharing Strategy. As illustrated in Fig. 2, Where2Comm [4] employs a confidence generator to estimate the likelihood of object presence in each BEV grid. Connected and autonomous vehicles (CAVs) then share regions where the ego has low confidence (*i.e.*, predicted as background) but neighboring CAVs are confident (*i.e.*, predicted as foreground). CoSDH [12] extends this idea with a supply–demand mechanism, using point cloud density (per pillar) as a proxy for observation reliability. Sparsely populated pillars are treated as uncertain and thus prioritized for CAVs’ support. While effective, both approaches largely disregard the informative cues embedded in background regions. CORE [18] indirectly addresses this by introducing a masked modeling strategy: it selects high-activated regions for transmission and applies a self-reconstruction objective at the receiver to recover the holistic scene. Although this implicitly captures some background context, it risks reconstructing irrelevant or noisy background content, potentially undermining perception robustness.

Efficient Communication Primitives. Beyond region-mask based sharing, several recent works explicitly design communication primitives to improve the bandwidth–accuracy trade-off. SlimComm [19] proposes Doppler-guided sparse queries to communicate a small set of motion-aware tokens for cooperative 3D perception. CoST [20] studies efficient collaboration from a unified spatiotemporal perspective, aiming to optimize information exchange jointly across time and agents. SparseAlign [21] develops a fully sparse framework for cooperative object detection, emphasizing sparse alignment and fusion to avoid dense BEV transmission.

III. PROBLEM FORMULATION

We study *3D object detection in collaborative perception* (CP) with N agents under the *intermediate fusion* paradigm. Each agent’s model comprises an encoder E_i , compressor ϕ_i , decompressor ψ_i , fusion module U_i , and decoder D_i , where $i \in \{1, 2, \dots, N\}$ indexes an agent.

Given the raw observation X_i , the encoder produces a BEV feature map:

$$F_i = E_i(X_i), \quad F_i \in \mathbb{R}^{C \times H \times W}, \quad (1)$$

where C is the channel dimension and (H, W) the BEV spatial resolution.

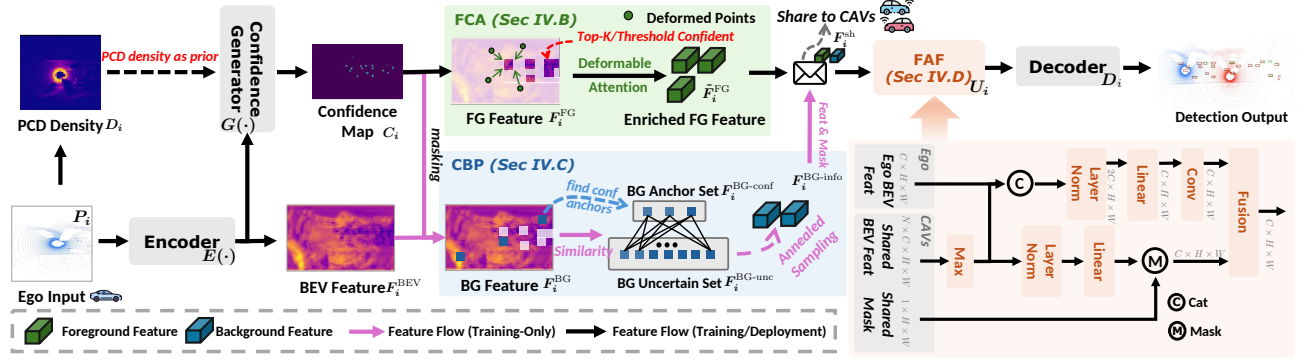


Fig. 3: Overview of FadeLead.

Since transmitting the full BEV feature map is infeasible, the compressor applies *sparse spatial sharing*:

$$\tilde{F}_i = \phi_i(F_i), \quad \tilde{F}_i \in \mathbb{R}^{C \times H \times W}. \quad (2)$$

This selects grid positions in the (H, W) plane and sends only their C -dimensional feature vectors. The sampled feature map \tilde{F}_i (extended to F_i^{sh} in our method), whose communication volume (in bits) is P_i , is then broadcast to neighbors within range δ .

When agent i receives \tilde{F}_j from a neighbor j with distance $d(i, j) \leq \delta$, it recovers the dense BEV feature map via the decompressor (e.g., scattering or reconstruction):

$$F_j = \psi_j(\tilde{F}_j), \quad F_j \in \mathbb{R}^{C \times H \times W}. \quad (3)$$

The fusion module U_i aggregates the ego feature with recovered neighbor features:

$$F_i^f = U_i(F_i, \{F_j \mid d(i, j) \leq \delta\}), \quad (4)$$

and the decoder outputs detections:

$$O_i = D_i(F_i^f). \quad (5)$$

Given ground-truth labels y_i , the training objective maximizes detection performance under a communication budget B :

$$\xi_{\Phi}(B) = \arg \max_{\theta} \sum_{i=1}^N g(O_i, y_i), \quad (6)$$

subject to

$$\sum_{j=1}^N |P_{j \rightarrow i}| \leq B, \quad (7)$$

where $|P_{j \rightarrow i}|$ denotes the communication volume of the message transmitted from agent j and received by agent i , and $g(\cdot, \cdot)$ is the evaluation metric (e.g., mAP).

IV. METHODOLOGY

A. Overview Framework

As illustrated in Fig. 3, each agent i begins by encoding its local point cloud \mathbf{P}_i with an encoder $E(\cdot)$ (e.g., PointPillars [22]) to produce an intermediate BEV feature map $\mathbf{F}_i^{\text{BEV}} \in \mathbb{R}^{C \times H \times W}$ and a density prior \mathbf{D}_i . A confidence generator $G(\cdot)$ estimates a confidence map $\mathbf{C}_i \in [0, 1]^{1 \times H \times W}$, which is used to partition the feature into foreground \mathbf{F}_i^{FG} and background \mathbf{F}_i^{BG} regions. These

regions are processed by two complementary branches: First, the *Foreground Context Attention* (FCA) branch enriches the foreground \mathbf{F}_i^{FG} by refining confidence with density priors and applying deformable attention over $\mathbf{F}_i^{\text{BEV}}$, yielding contextually enhanced features $\tilde{\mathbf{F}}_i^{\text{FG}}$. Second, the *Curricular Background Pruning* (CBP) branch (applied only during training) selects informative subsets $\mathbf{F}_i^{\text{BG-Info}}$ from the background and gradually anneals their transmission ratio. This curriculum encourages the model to encapsulate essential background context into the foreground representation.

During training, both $\tilde{\mathbf{F}}_i^{\text{FG}}$ and $\mathbf{F}_i^{\text{BG-Info}}$ are transmitted to collaborators, forming the shared feature \mathbf{F}_i^{sh} . At deployment, transmission is streamlined to include only the enriched foreground $\tilde{\mathbf{F}}_i^{\text{FG}}$, ensuring efficiency. Finally, each agent i collects shared features $\mathbf{F}_j^{\text{sh}} \mid j \in \mathcal{N}(i)$ from its neighbors and fuses them with its ego feature via the *Collaborative Context Fusion* (CCF) module. The fused representation is then decoded to produce the final detection outputs.

B. Foreground Context Attention (FCA)

Observation 1. Foreground-only sharing often suffers from false positives and incomplete semantics, as predicted foreground masks fail to account for contextual cues present in surrounding background regions.

Design 1. To address this, FCA enhances the reliability of foreground prediction in two ways.

First, unlike CoSDH [12], which uses point cloud (PCD) density only to measure ego demand for communication, we incorporate point cloud density as a prior to refine the predicted confidence map:

$$\mathbf{C}'_i = (1 - \text{norm}(\mathbf{D}_i)) \odot \mathbf{C}_i, \quad (8)$$

where \mathbf{C}_i is the confidence map, \mathbf{D}_i is the PCD density map, $\text{norm}(\cdot)$ denotes min-max normalization, and \odot indicates element-wise multiplication. This refinement primarily suppresses spurious false positives in low-density regions. To avoid over-suppressing true foregrounds that can also be high-density, we clamp the density penalty:

$$\mathbf{C}'_i = (1 - \min(\text{norm}(\mathbf{D}_i), \delta)) \odot \mathbf{C}_i, \quad (9)$$

where $\delta \in (0, 1)$ caps the maximum down-weighting (we use $\delta=0.5$ in experiments). Eq. (9) keeps dense foreground regions from being driven to near-zero confidence.

Second, the refined foreground features \mathbf{F}_i^{FG} are further enriched via multi-scale deformable attention [23] over the entire BEV map \mathbf{F}_i^{BEV} , producing context-aware foreground features $\tilde{\mathbf{F}}_i^{FG}$. By integrating context into the foreground, FCA mitigates semantic incompleteness and provides a more reliable foundation for subsequent feature sharing and fusion.

C. Curricular Background Pruning (CBP)

Observation 2. While background encodes valuable context (e.g., occlusion reasoning, scene layout), transmitting the entire background is infeasible. The central question is therefore how to exploit background cues during training while avoiding reliance on it at inference.

Design 2. As shown in Algorithm 1, CBP addresses this challenge by gradually shifting the model from background-assisted learning to foreground-only inference. At early training stages, CBP enriches learning with both foreground and informative background regions. Over time, it progressively reduces the background sharing ratio through a curriculum, forcing the model to internalize context into the shared foreground. Specifically, CBP consists of two key mechanisms:

Step 1: Informative background mining. Naively transmitting all background is wasteful, as most regions contain trivial empty space. To mitigate this, CBP partitions the background features \mathbf{F}_i^{BG} into two disjoint subsets: *confident background anchors* $\mathbf{F}_i^{BG-conf}$ (low-confidence, high-density areas that are almost certainly background), and *uncertain background* \mathbf{F}_i^{BG-unc} (regions where background assignment is less reliable due to sparsity, occlusion, or sensor noise). From the uncertain pool, CBP selects the most informative elements by measuring their similarity to confident background anchors. This procedure ensures that only the most representative background patterns are retained, so background is not discarded as noise but reinterpreted as a *scaffold* grounding foreground interpretation.

Step 2: Progressive pruning with a curriculum. To avoid over-reliance on background, CBP regulates its usage through a multi-stage annealing schedule. Training begins with a higher background sharing ratio r , analogous to a warm start, so the model benefits from abundant contextual cues. At each stage, this ratio is decayed by a factor γ , progressively pruning background features. The gradual reduction mimics a scaffolding process: early exposure to background provides support, while later pruning encourages the model to consolidate context within its foreground representation. By the final stage, background sharing is eliminated, and the model transmits only enriched foreground features that implicitly encode contextual priors.

D. Foreground Amplification Fusion (FAF)

Observation 3. Collaboration improves perception only if shared cues are fused effectively with the ego representation. Simple strategies such as concatenation or averaging are problematic: irrelevant or unaligned regions from collaborators introduce noise, and distributional mismatch between ego and remote features blurs salient activations. Without

Algorithm 1: Curricular Background Pruning (CBP)

Input: BEV features \mathbf{F}_i , density \mathbf{D}_i , confidence \mathbf{C}_i , foreground mask \mathbf{M}_i^{FG} , initial BG ratio r , similarity ratio τ , decay γ , image size $H \times W$

Output: Shared feature \mathbf{F}_i^{sh} , transmission mask \mathbf{M}_i^{sh}

- 1 **Notation.** $\text{TopK}_k(\cdot)$ selects the k elements with the highest score (right argument). $\text{sim}(\cdot, \cdot)$ denotes feature similarity (i.e., cosine).
 - 2 **Step 1: Informative Background Mining**
 - 3 **Function** $\text{MineBG}(\mathbf{F}_i, \mathbf{D}_i, \mathbf{C}_i, \mathbf{M}_i^{FG}, r, \tau)$:
 - 4 $\mathbf{M}_i^{BG} \leftarrow \mathbf{1} - \mathbf{M}_i^{FG}$ // background mask
 - 5 $\mathbf{C}_i^{BG} \leftarrow (\mathbf{1} - \mathbf{C}_i) \odot \mathbf{D}_i$ // refine confidence
 - 6 $k_a \leftarrow \lfloor r \cdot HW \rfloor, k_\tau \leftarrow \lfloor \tau \cdot HW \rfloor$
 - 7 $\mathbf{F}_i^{BG-conf} \leftarrow \text{TopK}_{k_a}(\mathbf{F}_i^{BG} | \mathbf{C}_i^{BG})$ // anchors
 - 8 $\mathbf{F}_i^{BG-unc} \leftarrow \mathbf{F}_i^{BG} \setminus \mathbf{F}_i^{BG-conf}$ // uncertain pool
 - 9 $\mathbf{F}_i^{BG-sel} \leftarrow \text{TopK}_{k_\tau}(\text{sim}(\mathbf{F}_i^{BG-unc}, \mathbf{F}_i^{BG-conf}))$
 - 10 $\mathbf{M}_i^{sh} \leftarrow \mathbf{M}_i^{FG} \cup \mathbf{M}_i^{BG-sel}$ // FG + mined BG
 - 11 **return** \mathbf{M}_i^{sh}
 - 12 **Step 2: Progressive Pruning (Curriculum)**
 - 13 **for** $\text{epoch} = 1, \dots, E$ **do**
 - 14 **for each training step / minibatch do**
 - 15 $\mathbf{M}_i^{sh} \leftarrow \text{MineBG}(\mathbf{F}_i, \mathbf{D}_i, \mathbf{C}_i, \mathbf{M}_i^{FG}, r, \tau)$
 - 16 $\mathbf{F}_i^{sh} \leftarrow \mathbf{F}_i \odot \mathbf{M}_i^{sh}$
 - 17 $r \leftarrow \gamma \cdot r$ // anneal BG ratio
 - 18 **return** $\mathbf{F}_i^{sh}, \mathbf{M}_i^{sh}$
-

careful design, fusion can degrade the feature map, making it less distinguishable rather than more informative.

Design 3. FAF addresses this by enforcing two principles: (1) fusion should explicitly model interactions between ego and received features, and (2) collaborative updates should amplify foreground cues while suppressing background noise.

Formally, let the ego feature be $\mathbf{F}_i^{\text{ego}} \in \mathbb{R}^{C \times H \times W}$, and let each neighbor $j \in \mathcal{N}(i)$ broadcast a shared feature \mathbf{F}_j^{sh} with mask \mathbf{M}_j^{sh} . At each timestep, ego i aggregates the received features by element-wise max fusion:

$$\mathbf{F}_i^r = \text{Proj}\left(\text{LN}\left(\max_{j \in \mathcal{N}(i)} \mathbf{F}_j^{sh}\right)\right), \quad \mathbf{M}_i^{sh} = \max_{j \in \mathcal{N}(i)} \mathbf{M}_j^{sh}.$$

Interaction Modeling. The normalized neighbor feature \mathbf{F}_i^r is combined with the ego representation to model cross-agent interactions:

$$\mathbf{F}_i^{\text{merge}} = \text{Conv}\left(\text{Proj}\left(\text{LN}\left([\mathbf{F}_i^{\text{ego}}, \mathbf{F}_i^r]\right)\right)\right).$$

Here, layer normalization equalizes activation scales, ensuring that foreground regions are not drowned out by noisy background responses.

Foreground Amplification. Finally, the merged feature is gated by the transmission mask and residually added to the ego feature:

$$\mathbf{F}_i^{\text{fused}} = \mathbf{F}_i^{\text{ego}} + \text{Proj}\left(\mathbf{F}_i^{\text{merge}} \odot \mathbf{F}_i^r \odot \mathbf{M}_i^{sh}\right).$$

This step selectively reinforces salient foreground activations while suppressing spurious background signals, ensuring

TABLE I: Detection accuracy and information selection ratio on OPV2V [13], V2X-R [24], and DAIR-V2X [25] datasets. “Ratio” denotes the proportion of BEV-plane information selected per collaborative agent. **Values** represent the best performance across all settings. **Values** indicate the performance gain when increasing the ratio from 1% to 10%, while **values** denote the improvement over the second-best method at the same ratio. * We adopt the intermediate fusion variant of CoSDH [12].

Method		OPV2V [13]		V2X-R [24]		DAIR-V2X [25]	
		AP@0.3/0.5/0.7↑	Ratio↓	AP@0.3/0.5/0.7↑	Ratio↓	AP@0.3/0.5/0.7↑	Ratio↓
Basic							
[-]	No Fusion	84.95/83.72/73.93	0%	72.52/71.09/61.06	0%	70.26/67.05/57.05	0%
[-]	Early Fusion	94.95/94.52/87.81	100%	88.17/87.82/77.73	100%	76.53/71.93/56.77	100%
[-]	Late Fusion	95.62/94.62/88.76	-	86.58/85.79/78.22	-	78.07/65.60/47.53	-
Dense Spatial Sharing							
[CVPR 2020]	When2Comm [5]	89.12/87.14/67.28	100%	86.01/84.74/74.75	100%	66.31/58.45/36.74	100%
[ECCV 2020]	V2VNet [16]	94.13/92.94/79.79	100%	85.12/83.65/67.92	100%	78.28/74.04/56.14	100%
[ECCV 2022]	V2XViT [17]	94.89/91.06/68.25	100%	91.17/88.93/78.19	100%	83.03/77.54/60.84	100%
[ICRA 2022]	AttFuse [13]	96.41 /94.79/81.05	100%	85.86/81.30/54.82	100%	82.99/78.72/62.02	100%
Sparse Spatial Sharing							
[NIPS 2022]	Where2Comm [4]	93.45/92.70/84.39	1%	90.24/89.09/77.55	1%	80.83/76.35/62.01	1%
		94.83/94.19/86.26	5%	91.74/90.47/79.37	5%	82.10/77.67/63.09	5%
		94.97/94.32/86.50	10%	92.00/90.68/79.76	10%	82.19/77.81/63.19	10%
		+1.52/+1.62/+2.11	1% → 10%	+1.76/+1.59/+2.21	1% → 10%	+1.36/+1.46/+1.18	1% → 10%
[ICCV 2023]	CORE [18]	51.52/50.46/40.72	1%	39.21/38.24/30.61	1%	44.48/42.27/32.59	1%
		80.96/79.59/67.18	5%	75.53/74.14/62.98	5%	59.83/56.16/44.60	5%
		86.10/84.57/71.19	10%	82.09/80.42/68.53	10%	64.26/60.09/47.59	10%
		+34.58/+34.11/+30.47	1% → 10%	+42.88/+42.18/+37.92	1% → 10%	+19.78/+17.82/+15.00	1% → 10%
[CVPR 2025]	CoSDH* [12]	90.54/89.06/77.70	1%	84.78/83.38/74.19	1%	78.89/74.40/61.30	1%
		92.40/90.99/79.61	5%	84.74/83.33/74.12	5%	82.10/77.89/64.02	5%
		93.89/92.74/81.67	10%	84.77/83.34/74.09	10%	82.72/78.57/64.35	10%
		+3.35/+3.68/+3.97	1% → 10%	-0.01/-0.04/-0.10	1% → 10%	+3.83/+4.17/+3.05	1% → 10%
[OURS]	FadeLead	95.81/95.05/88.10	1%	92.21/90.88/81.23	1%	83.27/79.10/64.45	1%
		(+2.36/+2.35/+3.71)		(+1.97/+1.79/+4.68)		(+2.44/+2.75/+2.44)	
		96.01/95.39/88.99	5%	92.90/91.71/83.25	5%	83.65/79.87/65.94	5%
		(+1.18/+1.07/+2.49)		(+1.16/+1.24/+3.88)		(+1.55/+1.98/+1.92)	
		96.01/ 95.41/89.02	10%	92.97/91.79/83.32	10%	83.52/79.76/65.98	10%
		(+1.04/+1.09/+2.52)		(+0.97/+1.11/+3.56)		(+0.8/+1.19/+1.63)	
		+0.20/+0.36/+0.92	1% → 10%	+0.06/+0.91/+2.09	1% → 10%	+0.25/+0.66/+1.53	1% → 10%

Bandwidth consumption (100% for dense, 1/5/10% for sparse): V2VNet = V2XViT = AttFuse > When2Comm > Early > realworld bandwidth constraint 28Mbps [26] > Where2Comm = CORE > CoSDH = FadeLead > Late. Though bandwidth-efficient, late fusion inevitably incurs additional latency (e.g., full model inference, post-processing).

that collaboration strengthens the representation rather than injecting noise. The result is a cleaner, more discriminative fused feature (visualized in Fig. 5 later).

V. EVALUATION

A. Experimental Setup

Datasets. We evaluate our method on both **simulated** and **real-world** datasets (OPV2V [13], V2X-R [24] and DAIR-V2X [25]) for LiDAR-based 3D object detection. These datasets span varying numbers of collaborative agents, object density, and lane conditions.

Baselines. We compare our method against recent open-source approaches, including (1) naive settings, (2) dense spatial sharing methods that transmit the full intermediate BEV feature (i.e., When2Comm [5], V2VNet, V2XViT [17], and AttFuse [13]) and (3) sparse sharing methods that follow a foreground-centric strategy (i.e., Where2Comm [4], CORE [18], and CoSDH [12]). For CoSDH, which employs an intermediate-late hybrid fusion strategy, we focus on the intermediate fusion variant, as late fusion is less practical for real-world deployment (especially in high-speed scenarios) due to the additional latency introduced by complete model inference and post-processing (e.g., NMS).

Implementation. For fair comparison, we standardize the spatial coverage and resolution of intermediate BEV features. Specifically, OPV2V and V2X-R adopt a perception range of [-140.8, -38.4, -3, 140.8, 38.4, 1] with a downsampling rate of 4, while DAIR-V2X uses [-100.8, -40, -3.5, 100.8, 40, 1.5] with a downsampling rate of 2. All baseline methods are retrained with their original hyperparameters and use PointPillars [22] with voxel size [0.4, 0.4, h]m as the encoder. Under these settings, the resulting BEV feature map resolutions are 176×48 for OPV2V/V2X-R and 126×50 for DAIR-V2X. For CoSDH and FadeLead, we apply channel compression that downsamples the shared feature map from $[H, W, C]$ to $[H, W, C/R]$ ($C=256$ and $R=16$ in our settings) along with quantization (float32→float16), which further reduces bandwidth consumption. To better reflect real-world deployment constraints, we avoid projecting point clouds into the ego coordinate during preprocessing (a step adopted in some prior works for training stability, but impractical before feature sharing in real deployments). For CBP, we set ($r=0.1$, $\gamma=0.8$) with decay every 5 epochs. Full hyperparameters and configurations are provided in anonymously released code. All models are trained on four Nvidia RTX 3090 GPUs.

Evaluation. We evaluate average precision (AP) at IoU

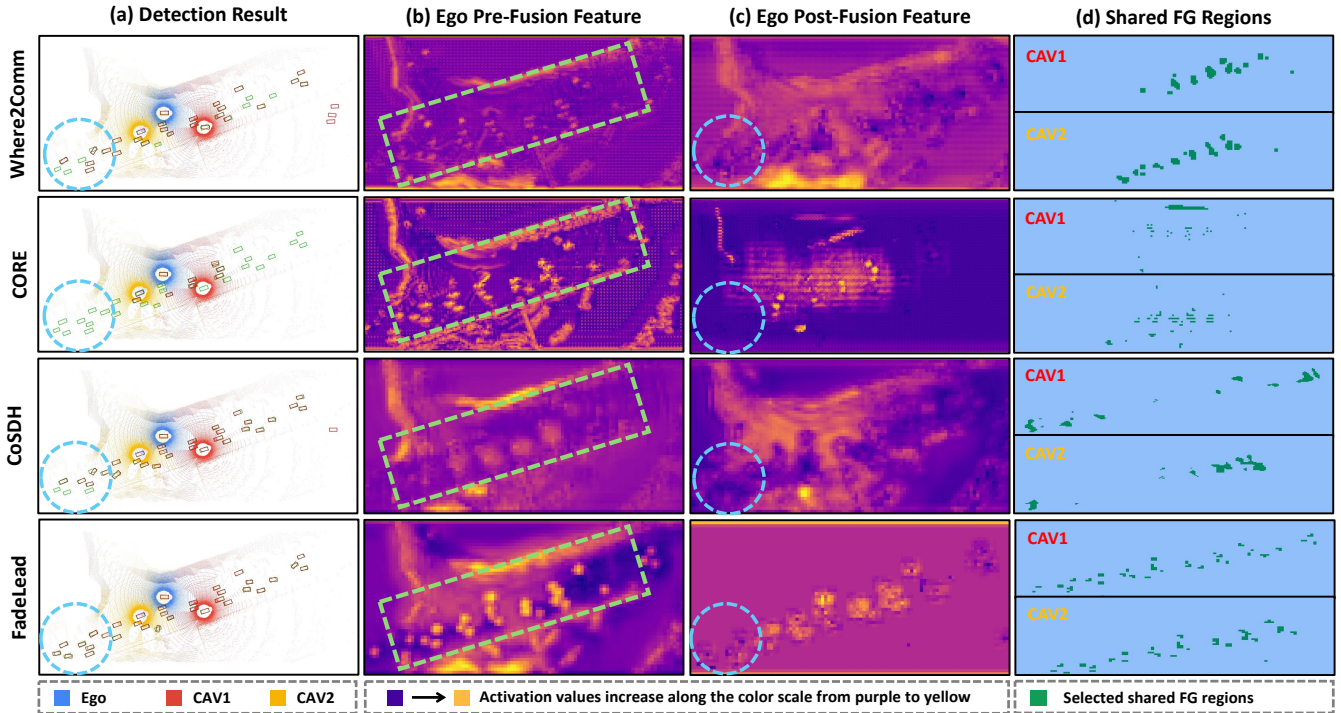


Fig. 4: Visualization on OPV2V [13] with top-1% confident foreground selection. We show (a) detection results, (b) the ego BEV feature before fusion, (c) the ego BEV feature after fusion, and (d) the shared regions from collaborating CAVs.

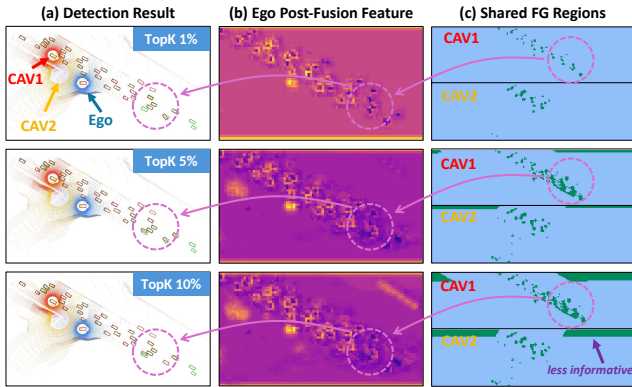


Fig. 5: Visualization of detection results, post-fusion features, and shared FG regions under different top- k ratios on V2X-R [24].

thresholds of 0.3, 0.5 and 0.7 for 3D object detection. To examine the effectiveness of spatial information selection, we vary the selection ratio to 1%, 5%, and 10%. These values are chosen as the foreground typically occupies $\leq 10\%$ of the BEV plane across all scenarios, and the resulting bandwidth remains within the limits of vehicular networks [26].

B. Quantitative Results

Better performance even under extreme low network bandwidth. We evaluate FadeLead against baseline methods using AP@0.3/0.5/0.7 under varying information selection ratios. As shown in Table I, dense spatial sharing methods (*i.e.*, When2Comm [5], V2VNet [16], V2XViT [17], AttFuse [13]) can achieve strong detection accuracy but only by transmitting the entire BEV feature, making them bandwidth-inefficient for practical deployment. In contrast, sparse sharing approaches transmit only a small fraction of informative regions while often achieving comparable or

superior accuracy at significantly reduced bandwidth cost. Within the sparse sharing category, FadeLead consistently outperforms SOTA sparse spatial sharing methods across all datasets at corresponding selection ratios, as indicated by the values in Table I. At the most restrictive 1% selection ratio, FadeLead demonstrates substantial improvements: +2.36/+2.35/+3.71 AP on OPV2V, +1.97/+1.79/+4.68 AP on V2X-R, and +2.44/+2.75/+2.44 AP on DAIR-V2X compared to the second-best method. This low bandwidth requirement is particularly advantageous in dense collaborative scenarios (*e.g.*, intersections with multiple participating vehicles) and provides better resilience in unstable vehicular networks.

Background contextual cues are effectively encapsulated.

The values in Table I show that FadeLead reaches near-optimal performance even at the 1% selection ratio, with only marginal improvements when increasing the ratio to 10% (*e.g.*, +0.20/+0.36/+0.92 on OPV2V). This small gap indicates that essential background context is already encapsulated into the minimal set of shared foreground regions, leaving little room for further gains by transmitting more background. In contrast, reconstruction-based method CORE exhibits severe instability, with large fluctuations across ratios (*e.g.*, +34.58/+34.11/+30.47 on OPV2V). The stability of FadeLead demonstrates that it integrates contextual cues efficiently at extremely low bandwidth, providing both robustness and scalability for real-world deployment.

FadeLead with improved bandwidth. FadeLead is primarily designed for extreme bandwidth constraints, where only predicted foreground region features are transmitted. A natural question, however, is whether additional bandwidth can further enhance its performance. As shown in Fig. 5c, increasing the top- k ratio first brings in small portions

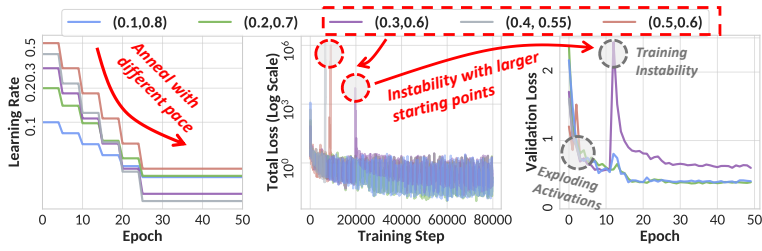


Fig. 6: Effect of curriculum pacing under different initialization points and progression rates on OPV2V [13].

of background adjacent to the foreground (*i.e.*, 1%→5%), which provide useful complementary context to correct feature maps. As the ratio grows further (*i.e.*, 5%→10%), more certain background regions (typically near scene boundaries) are included. While less informative, these regions can still offer minor benefits for refining detection results.

Pace sensitivity of curriculum. We evaluate the sensitivity of the curriculum schedule by varying the (initial background ratio r , the decay factor γ). All schedules are designed to decay at comparable rates and eventually converge to a near-zero background ratio. As illustrated in Fig. 6, larger initial ratios induce training instabilities (*e.g.*, sudden spikes in validation loss, numerical divergence), whereas smaller ratios lead to stable convergence. These results suggest that introducing a modest amount of background early on provides useful context to stabilize learning, but excessive background may overshadow foreground and disrupt training.

C. Qualitative Results

FadeLead enhances foreground representations. Figure 4b shows that our FCA module strengthens the semantic quality of the ego’s pre-fusion features. Compared with Where2Comm [4] and CoSDH [12], FadeLead produces sharper activations at target object locations, reflecting better semantic alignment. Unlike CORE [18], which suppresses background responses uniformly, FadeLead retains higher activation in nearby background regions where the ego has strong observability. Though these are background areas, maintaining confident activation is beneficial: it encodes the ego’s reliable observations and helps distinguish certain background from uncertain or occluded regions. This not only improves pre-fusion representations but also provides a more trustworthy input to the subsequent collaborative fusion stage, ensuring that the shared and fused features are grounded in clear observability semantics.

FadeLead enables more reliable feature sharing and fusion. Figure 4 highlights how different collaborative perception methods vary in foreground prediction accuracy, feature-map quality after fusion, and the reliability of shared regions. For foreground estimation and sharing, as shown in Fig. 4d, Where2Comm often misses distant objects due to inaccurate foreground prediction, leading to incomplete sharing and weak fusion. CORE’s reconstruction-based design relies heavily on abundant source information, and under limited input, it can even corrupt the ego’s pre-fusion feature map. CoSDH leverages PCD density as demanding, but many of its selected regions remain ambiguous to the

TABLE II: Ablation study.

			OPV2V [13]			DAIR-V2X [25]		
FCA	CBP	FAF	AP@0.3	AP@0.5	AP@0.7	AP@0.3	AP@0.5	AP@0.7
×/✓	×	×	collapse	collapse	collapse	collapse	collapse	collapse
✓	×	✓	collapse	collapse	collapse	80.45	76.15	61.84
						(-3.07)	(-3.61)	(-4.14)
×	✓	✓	96.22	95.17	87.12	83.24	79.19	64.91
			(-0.21)	(-0.24)	(-1.9)	(-0.28)	(-0.57)	(-1.07)
✓	✓	×	96.30	94.19	84.13	79.82	73.54	57.83
			(+0.29)	(-1.22)	(-4.89)	(-3.7)	(-6.22)	(-8.15)
✓	✓	✓	96.01	95.41	89.02	83.52	79.76	65.98

TABLE III: Performance of intermediate-late fusion variant.

Method	OPV2V			V2X-R			DAIR-V2X		
	AP@0.3	AP@0.5	AP@0.7	AP@0.3	AP@0.5	AP@0.7	AP@0.3	AP@0.5	AP@0.7
Where2Comm	96.43	96.01	90.40	93.60	92.55	83.08	83.26	78.29	64.03
CORE	95.41	94.67	87.62	90.81	89.64	80.73	62.14	58.32	46.82
CoSDH	97.36	96.75	90.79	92.39	91.72	85.03	84.59	79.99	66.01
FadeLead	96.82	96.48	92.51	93.56	92.77	86.89	83.98	79.66	66.51

ego, resulting in weak post-fusion activations. In contrast, FadeLead combines FCA’s explicit supervision with PCD density-smoothed confidence to produce more reliable region selection, and with FAF reinforcing these cues after fusion, its post-fusion feature map (Fig. 4c) becomes more clearly distinguishable between background and foreground.

VI. ABLATION & DISCUSSION

Ablation study. Table II validates the effect of each component in FadeLead. Without CBP, strictly sharing predicted foreground features is prone to collapse during early training. Even when convergence is achieved, excluding CBP causes notable performance drop on DAIR-V2X, as the model loses the gradual curriculum that both regularizes learning and injects background context into the foreground representation. Removing FCA degrade performance especially at higher IoUs, indicating that FCA enhances precise localization by aggregating contextual cues around foreground objects (see Fig. 5). Disabling FAF yields the sharpest decrease, particularly on DAIR-V2X, underscoring its role in effectively fusing local and shared features.

Hybrid with late fusion. While intermediate-late hybrid fusion inevitably incurs additional latency (*e.g.*, full model inference, post-processing) and is therefore less practical for real-world deployment, Table III shows that our method achieves comparable performance to the state-of-the-art CoSDH [12] at AP@0.3 and AP@0.5, and consistently outperforms it at higher IoU thresholds (*i.e.*, AP@0.7).

Further results. Due to limited space, additional experimental details at top-5%/10% will be provided as supplement.

VII. CONCLUSION

We presented FadeLead, a foreground-centric collaborative perception framework that enriches transmitted foreground region features with context to enable efficient and robust collaborative perception. Through a curricular training strategy that internalizes background context into foreground representations and fusion mechanisms that emphasize salient information, FadeLead push the limits of foreground-centric sharing. Extensive experiments on both simulation and real-world datasets validate its effectiveness under stringent communication constraints, underscoring its efficiency for deployment in practical vehicular networks.

REFERENCES

- [1] X. Gao, R. Xu, J. Li, Z. Wang, Z. Fan, and Z. Tu, “STAMP: Scalable Task And Model-agnostic Collaborative Perception,” Jan. 2025.
- [2] Y. Lu, Y. Hu, Y. Zhong, D. Wang, Y. Wang, and S. Chen, “An Extensible Framework for Open Heterogeneous Collaborative Perception,” Apr. 2024.
- [3] X. Gao, Y. Wu, R. Wang, C. Liu, Y. Zhou, and Z. Tu, “Langcoop: Collaborative driving with language,” in *Proceedings of the Computer Vision and Pattern Recognition Conference, 2025*, pp. 4226–4237.
- [4] Y. Hu, S. Fang, Z. Lei, Y. Zhong, and S. Chen, “Where2comm: Communication-Efficient Collaborative Perception via Spatial Confidence Maps,” Sept. 2022.
- [5] Y.-C. Liu, J. Tian, N. Glaser, and Z. Kira, “When2com: Multi-Agent Perception via Communication Graph Grouping,” June 2020.
- [6] D. Yang, K. Yang, Y. Wang, J. Liu, Z. Xu, R. Yin, P. Zhai, and L. Zhang, “How2comm: Communication-Efficient and Collaboration-Pragmatic Multi-Agent Perception,” 2023.
- [7] Y. Hu, J. Peng, S. Liu, J. Ge, S. Liu, and S. Chen, “Communication-efficient collaborative perception via information filling with codebook,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024*, pp. 15 481–15 490.
- [8] Y. Lu, Q. Li, B. Liu, M. Dianati, C. Feng, S. Chen, and Y. Wang, “Robust collaborative 3d object detection in presence of pose errors,” *arXiv preprint arXiv:2211.07214*, 2022.
- [9] X. Zhang, Z. Zhou, Z. Wang, Y. Ji, Y. Huang, and H. Chen, “Co-mtp: A cooperative trajectory prediction framework with multi-temporal fusion for autonomous driving,” *arXiv preprint arXiv:2502.16589*, 2025.
- [10] K. Yang, D. Yang, K. Li, D. Xiao, Z. Shao, P. Sun, and L. Song, “Align before collaborate: Mitigating feature misalignment for robust multi-agent perception,” in *European Conference on Computer Vision*. Springer, 2024, pp. 282–299.
- [11] S. Huang, J. Zhang, Y. Li, and C. Feng, “ActFormer: Scalable Collaborative Perception via Active Queries,” Mar. 2024.
- [12] J. Xu, Y. Zhang, Z. Cai, and D. Huang, “Cosdh: Communication-efficient collaborative perception via supply-demand awareness and intermediate-late hybridization,” *arXiv preprint arXiv:2503.03430*, 2025.
- [13] R. Xu, H. Xiang, X. Xia, X. Han, J. Li, and J. Ma, “Opv2v: An open benchmark dataset and fusion pipeline for perception with vehicle-to-vehicle communication,” in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, p. 2583–2589.
- [14] Q. Chen, S. Tang, Q. Yang, and S. Fu, “Cooper: Cooperative Perception for Connected Autonomous Vehicles based on 3D Point Clouds,” May 2019.
- [15] E. Arnold, M. Dianati, R. De Temple, and S. Fallah, “Cooperative perception for 3d object detection in driving scenarios using infrastructure sensors,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 3, pp. 1852–1864, 2020.
- [16] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, J. Tu, and R. Urtasun, “V2VNet: Vehicle-to-Vehicle Communication for Joint Perception and Prediction,” Aug. 2020.
- [17] R. Xu, H. Xiang, Z. Tu, X. Xia, M.-H. Yang, and J. Ma, “V2X-ViT: Vehicle-to-Everything Cooperative Perception with Vision Transformer,” Aug. 2022.
- [18] B. Wang, L. Zhang, Z. Wang, Y. Zhao, and T. Zhou, “Core: Cooperative Reconstruction for Multi-Agent Perception,” in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*. Paris, France: IEEE, Oct. 2023, pp. 8676–8686.
- [19] M. Yazgan, Q. Wu, I. Hamdard, S. Li, and J. M. Zoellner, “Slimcomm: Doppler-guided sparse queries for bandwidth-efficient cooperative 3-d perception,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision, 2025*, pp. 1782–1791.
- [20] Z. Tang, Y. Liu, Y. Sun, Y. Gao, J. Chen, R. Xu, and S. Liu, “Cost: Efficient collaborative perception from unified spatiotemporal perspective,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision, 2025*, pp. 1120–1129.
- [21] Y. Yuan, Y. Xia, D. Cremers, and M. Sester, “Sparsealign: A fully sparse framework for cooperative object detection,” in *Proceedings of the Computer Vision and Pattern Recognition Conference, 2025*, pp. 22 296–22 305.
- [22] A. H. Lang, S. Vora, H. Caesar, L. Zhou, J. Yang, and O. Beijbom, “Pointpillars: Fast encoders for object detection from point clouds,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, 2019*, pp. 12 697–12 705.
- [23] X. Zhu, W. Su, L. Lu, B. Li, X. Wang, and J. Dai, “Deformable detr: Deformable transformers for end-to-end object detection,” *arXiv preprint arXiv:2010.04159*, 2020.
- [24] X. Huang, J. Wang, Q. Xia, S. Chen, B. Yang, C. Wang, and C. Wen, “V2x-r: Cooperative lidar-4d radar fusion for 3d object detection with denoising diffusion,” *arXiv preprint arXiv:2411.08402*, 2024.
- [25] H. Yu, Y. Luo, M. Shu, Y. Huo, Z. Yang, Y. Shi, Z. Guo, H. Li, X. Hu, J. Yuan, and Z. Nie, “Dairv2x: A large-scale dataset for vehicle-infrastructure cooperative 3d object detection,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022*, pp. 21 361–21 370.
- [26] R. Bera, J. Bera, S. Sil, S. Dogra, N. Sinha, and D. Mondal, “Dedicated short range communications (dsrc) for intelligent transport system,” in *2006 IFIP International Conference on Wireless and Optical Communications Networks*, 2006, pp. 5 pp.–5.