

ModalPatch: A Plug-and-Play Module for Robust Multi-Modal 3D Object Detection under Modality Drop

Shuangzhi Li¹, Lei Ma^{2,1}, and Xingyu Li¹

¹University of Alberta, Canada ²The University of Tokyo, Japan
 {shuangzh, xingyu}@ualberta.ca, ma.lei@acm.org

Abstract—Multi-modal 3D object detection is pivotal for autonomous driving, integrating complementary sensors like LiDAR and cameras. However, its real-world reliability is challenged by transient data interruptions and missing, where modalities can momentarily drop due to hardware glitches, adverse weather, or occlusions. This poses a critical risk, especially during a simultaneous modality drop, where the vehicle is momentarily blind. To address this problem, we introduce ModalPatch, the first plug-and-play module designed to enable robust detection under arbitrary modality-drop scenarios. Without requiring architectural changes or retraining, ModalPatch can be seamlessly integrated into diverse detection frameworks. Technically, ModalPatch leverages the temporal nature of sensor data for perceptual continuity, using a history-based module to predict and compensate for transiently unavailable features. To improve the fidelity of the predicted features, we further introduce an uncertainty-guided cross-modality fusion strategy that dynamically estimates the reliability of compensated features, suppressing biased signals while reinforcing informative ones. Extensive experiments show that ModalPatch consistently enhances both robustness and accuracy of state-of-the-art 3D object detectors under diverse modality-drop conditions. Code will be available at https://github.com/Castiel-Lee/MM3Det_MD.

I. INTRODUCTION

Multi-modal 3D object detection [1], [2] has become a cornerstone in applications such as autonomous driving and robotics. By integrating complementary sensing modalities such as LiDAR and cameras, these systems achieve more accurate and reliable object localization and classification. However, despite these benefits, multi-modal systems face critical robustness challenges during real-world deployment. Sensor inputs are often susceptible to failure [3], [4] due to a variety of unpredictable conditions, including hardware malfunction [5], [6], adverse weather [7], [4], and occlusion [8], [6]. Moreover, asynchronous sensors may produce misaligned or incomplete data because of differences in sampling frequencies [9]. These issues result in modality drop (i.e., temporary or partial loss of sensor inputs) that degrade detection performance if not properly handled. All of these real-world challenges point to a fundamental problem in multi-modal 3D object detection:

How can we ensure robust detection when one or more modalities drop unexpectedly?

Existing approaches to this problem [10], [11], [12] have two significant limitations. First, they primarily address a simplified scenario: *dependent* modality-drop, where at least




Fig. 1. (a) The proposed ModalPatch framework compensates for arbitrary modality drops (LiDAR or camera) and can be seamlessly integrated into existing detectors without retraining them. (b) Performance boost achieved by ModalPatch under the 30% modality drop rate for various detectors.

one modality (e.g., LiDAR or camera) is guaranteed to remain available. This overlooks a more critical and realistic scenario: a simultaneous modality drop, where transient issues cause all sensors to lose signal concurrently for short periods. Though brief, such events create moments of total perceptual blindness for an autonomous system. Second, these solutions often require re-designing the core detection architecture or complete model retraining [11], [12], making them resource-intensive, inflexible, and difficult to generalize across different state-of-the-art (SOTA) detectors.

To bridge these research gaps, we propose ModalPatch, a lightweight, plug-and-play solution designed to make multi-modal 3D detectors resilient to transient modality drops, regardless of whether the drop is dependent or simultaneous (independent). It can be seamlessly integrated into a variety of existing detection frameworks without requiring retraining or architecture modifications. Specifically, our key insight is that 3D perception systems in domains like autonomous driving and robotics typically operate on temporal data streams, where sensory inputs evolve continuously over time. ModalPatch exploits this property by maintaining a short-term historical memory of past feature representations, which it uses to predict and compensate for current missing modality data. To further improve robustness, ModalPatch

incorporates a cross-modality fusion module that promotes interaction between modalities, allowing the system to mitigate the weaknesses of individual sensors. However, predictions based on temporal history may carry noise or bias. To address this, we introduce an uncertainty-aware mechanism in our cross-modality fusion module. Estimating the reliability of both the predicted and live features, this mechanism suppresses unreliable components and amplifies trustworthy ones, yielding more robust detection outcomes. We validate ModalPatch through extensive experiments on multiple SOTA 3D detectors, demonstrating its ability to generalize across architectures and consistently improve performance under diverse modality-drop conditions.

Our contributions are summarized as follows:

- We propose ModalPatch, the first plug-and-play solution that can be seamlessly integrated into existing 3D object detection frameworks to handle arbitrary modality drop.
- We leverage the temporal nature of 3D perception by using historical feature memory to predict missing modality features. This serves as an effective and adaptive compensation mechanism in dynamic environments where sensor inputs may drop unexpectedly.
- We introduce an uncertainty-guided cross-modality fusion strategy. This mechanism suppresses unreliable signals and reinforces trustworthy ones, improving robustness against both feature bias and noise propagation.
- We demonstrate that ModalPatch consistently improves robustness on various SOTA detectors and generalizes well across different 3D detection architectures and a wide range of modality-drop scenarios.

II. RELATED WORKS

A. Multi-modal 3D Object Detection

Multi-modal 3D object detection [1], [2], [13] has rapidly advanced by leveraging the strengths of multiple modalities (e.g., LiDAR and camera). Based on feature extraction from heterogeneous modalities, recent methods can be broadly categorized into two main branches: 1) BEV-based feature fusion, where features from different modalities are projected into a shared bird’s-eye-view (BEV) representation to facilitate cross-modal alignment and joint reasoning, as exemplified by BEVFusion [14], [15] and its extensions [11], [16]; 2) transformer-based cross-modal interaction, where heterogeneous modality features are fused with attention mechanisms via spatial-aware queries with unified 3D position embedding [17], as in Cross-Modal Transformer (CMT) [10], TransFusion [18], and SparseLIF [19]. While these detectors achieve strong performance under normal conditions, they often experience significant degradation when one or more modalities are missing.

B. Multi-modal 3D Object Detection with Modality Drop

Traditional trials adopt the Kalman filter to compensate for missing inputs by estimating the target’s temporal state [20]. Recent studies have explored feature compensation for multi-modal 3D object detection: CMT [10] enhances robustness by randomly masking camera inputs during training;

UniBEV [11] introduces unified BEV queries for feature extraction and channel-adaptive fusion; and MEFFormer [12] employs dynamic ensembling of detectors trained on different modality subsets to handle both single- and multi-modality inputs. While these approaches represent important progress, they still fall short in addressing the challenges highlighted earlier: they primarily assume dependent modality drop (at least one modality exists) rather than truly independent random drop (no modality input at some time points), and they often require substantial architectural re-design or retraining, limiting their practicality. In contrast, our work introduces a plug-and-play module that can be seamlessly integrated into existing detectors to handle arbitrary modality drops, thereby ensuring robust performance without sacrificing flexibility or generality.

III. METHODOLOGY

ModalPatch is a lightweight “patch” module that can be seamlessly plugged into diverse detection frameworks for various modality drop scenarios. As demonstrated in Fig. 1, by simply attaching this module, various 3D detectors can be patched to recover missing modalities — whether single or multiple, thus maintaining robust detection performance in the presence of arbitrary modality loss. As shown in Fig. 2, ModalPatch combines (1) *history-based feature prediction*, which leverages temporal continuity to compensate for missing features, and (2) *uncertainty-guided cross-modality fusion*, which improves these features through uncertainty-aware cross-domain complementary refinement.

A. History-based Feature Prediction (HFP)

Temporal continuity is inherent in multi-modal 3D object detection and has been extensively explored in recent studies [16], [21]. Inspired by [21], we design a history-based feature prediction module that leverages temporal observations to compensate for missing modalities. Specifically, we model the evolution of features over a sequence of past frames and use the most recent historical memory to predict the current features, which are then employed as substitutes for the missing modality inputs. Moreover, during deployment, the compensated features can serve as an alternative history to ensure continuous compensation and provide relatively stable detection performance over time.

History-based Temporal Transformer. The input images and point clouds are first processed by their respective backbones to obtain feature maps $F_M^{t_0} \in \mathbf{R}^{B \times D_M \times H_M \times W_M}$, where t_0 denotes the current time and $\{B, D_M, H_M, W_M\}$ denote the batch size, channel dimension, height, and width of the extracted features for modality $M \in \{\text{img}, \text{pts}\}$. As shown in Fig. 3, we maintain a memory bank of historical features $\{F_M^{t_0-1}, F_M^{t_0-2}, \dots, F_M^{t_0-\tau}\}$ with a temporal length of τ for each modality M to store historical information. Then we concatenate these history features as S_M and use them as the key and value, while the learnable BEV embedding Q_M serves as the query to model feature dynamics across time. Specifically, we adopt spatially sensitive deformable attention [22] to adaptively aggregate local context

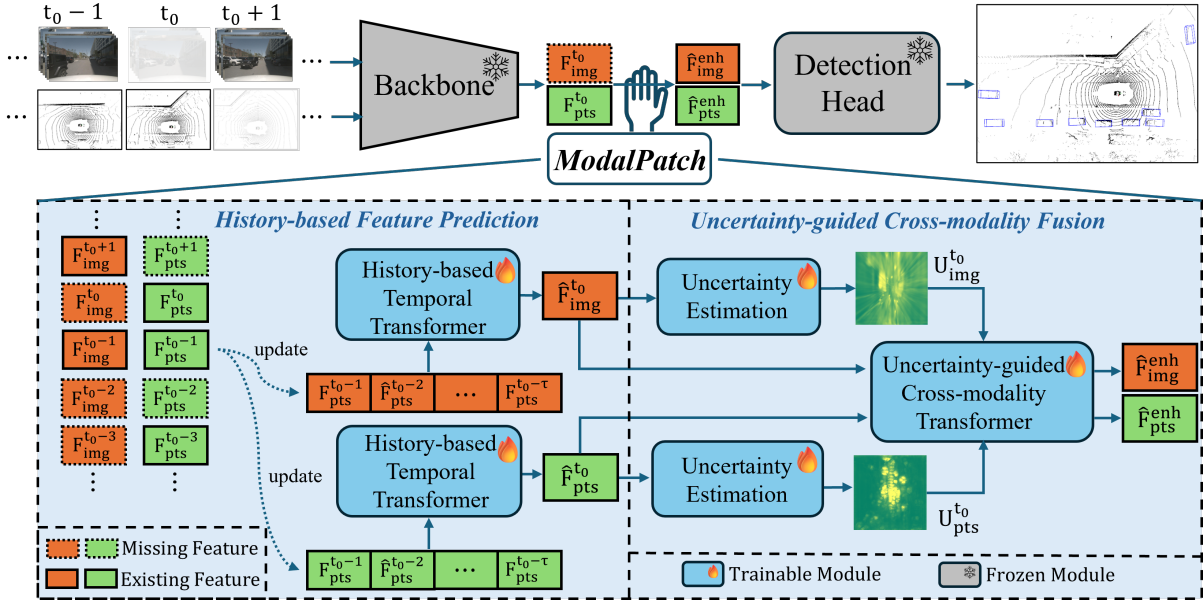


Fig. 2. Overview of the proposed ModalPatch module. Given multi-modal inputs (e.g., point clouds by LiDAR and images by camera), features are first extracted by the frozen backbone. To address possible modality-drop scenarios, the plug-and-play ModalPatch introduces two modules: (1) a history-based feature prediction module, which leverages temporal information from past frames to predict current features and provide initial compensation; and (2) an uncertainty-based cross-modality fusion module, which estimates spatial uncertainty and employs cross-modality complementary information to enhance compensated features.

and capture fine-grained temporal dynamics:

$$\tilde{F}_M^{t_0} = \text{DeformAttn}(\text{DeformAttn}(Q_M, S_M), S_M), \quad (1)$$

where $\tilde{F}_M^{t_0}$ denotes the compensated temporal dynamics. We then add this dynamic to the most recent historical feature $F_M^{t_0-1}$ to obtain the compensated feature for the current time:

$$\hat{F}_M^{t_0} = F_M^{t_0-1} + \tilde{F}_M^{t_0}. \quad (2)$$

Since training datasets collected in autonomous driving and robotics generally provide complete multi-modal inputs, we reasonably assume that modality missing does not occur during training. Thus, the ground-truth feature $F_M^{t_0}$ can be used to supervise prediction accuracy of $\hat{F}_M^{t_0}$ via an ℓ_2 loss [23]:

$$\mathcal{L}_{\text{TempPred}} = \|\hat{F}_M^{t_0} - F_M^{t_0}\|_2^2. \quad (3)$$

History Memory Update. During training, the history memory bank is updated using the ground-truth features $F_M^{t_0-1}$ extracted by the backbone, which ensures accurate temporal supervision. At inference time, for modalities that are available, the bank is updated with the extracted features $F_M^{t_0-1}$, while for missing modalities, we instead use the compensated features $\hat{F}_M^{t_0-1}$ as substitutes. This strategy guarantees temporal continuity of the memory bank, allowing the model to maintain reliable predictions over consecutive time steps even under arbitrary modality drop.

B. Uncertainty-guided Cross-modality Fusion (UCF)

Although modality compensation alleviates the problem of missing inputs, each single modality still suffers from inherent limitations [1], [14], such as the restricted field of view and incomplete geometrical cues. To overcome these

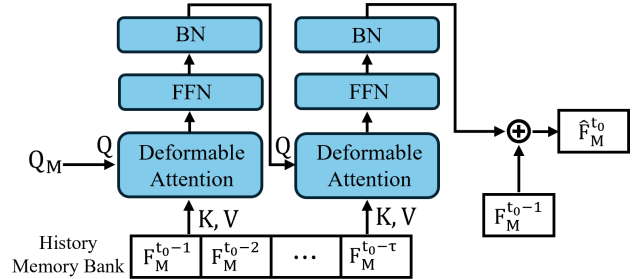


Fig. 3. History-based temporal transformer, taking learnable queries and the history memory bank as inputs to generate compensated features.

constraints, we perform cross-modality fusion to leverage complementary information across sensors. However, due to error accumulation in temporal prediction and modality-specific deficiencies mentioned above, the compensated features may contain unignorable biases, which can propagate through the fusion process and degrade detection performance. To mitigate this issue, we design an uncertainty-guided fusion module, which explicitly estimates the reliability of compensated features and adaptively weights their contributions. This mechanism suppresses biased or unreliable signals while reinforcing trustworthy ones, thereby enabling robust cross-modality feature enhancement and improving overall detection accuracy.

Uncertainty Estimation. Uncertainty estimation has been widely studied as an effective approach to assess the reliability of model behavior [24], [25]. To mitigate potential cumulative errors and bias in compensated features, we propose to estimate the spatial uncertainty of the compensated features and use it as a reliability measure to adaptively weight their

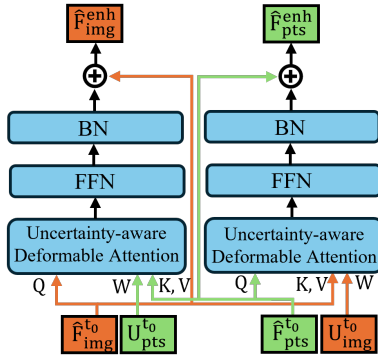


Fig. 4. Uncertainty-guided cross-modality transformer, taking compensated features and uncertainty maps of two modalities as inputs to cross-modally enhance features.

contributions in subsequent cross-modality fusion.

Concretely, for each spatial location in the compensated feature map $\hat{F}_M^{t_0}$, we assume the predicted value follows a Gaussian distribution centered at the corresponding ground-truth feature $F_M^{t_0}$, with an unknown variance σ_M^2 . To estimate this uncertainty, we employ a lightweight MLP to regress the variance from $\hat{F}_M^{t_0}$, producing a variance map:

$$\sigma_M^2 = \text{MLP}(\hat{F}_M^{t_0}) \in \mathbf{R}^{B \times 1 \times H_M \times W_M}. \quad (4)$$

To optimize the MLP, we adopt the negative log-likelihood (NLL) loss [26], which encourages the predicted variances to reflect the reliability of the compensation.

$$\mathcal{L}_{\text{Uncert}} = \sum_{i=1}^{H_M W_M} \frac{1}{2} \left(\frac{\|\hat{F}_{M,i}^{t_0} - F_{M,i}^{t_0}\|_2^2}{\sigma_{M,i}^2} + \log \sigma_{M,i}^2 + \log(2\pi) \right), \quad (5)$$

where $\sigma_{M,i}^2$ denotes the predicted variance at spatial location i . Minimizing this loss encourages the model to assign higher variance to less reliable regions, thereby allowing σ_M^2 to serve as an uncertainty estimate. We then define the uncertainty map $U_M^{t_0} := \sigma_M$, which captures the spatial reliability of the compensated features $\hat{F}_M^{t_0}$ and guide the subsequent cross-modality fusion to suppress uncertain regions while emphasizing more confident signals.

Cross-modality Fusion. Cross-modality fusion is designed to address the potential unreliability of compensated features, which may contain biases or noise from temporal prediction as well as inherent limitations of individual modalities. By leveraging complementary information across modalities, the fusion process can suppress unreliable signals, recover missing details, and thereby enhance the robustness of the compensated features. Here, we innovate an uncertainty-aware deformable transformer block to transfer reliable complementary information from one modality to the other. Notably, on the basis of deformable attention [22], we incorporate the uncertainty measurement of the compensated features into attention calculation.

Let's take the enhancement of the image feature by the LiDAR feature as an example. As shown in Fig. 4, to ensure sensitivity to the spatial distribution of images, we use the compensated image feature $\hat{F}_{img}^{t_0}$ as the query, while the

compensated LiDAR feature $\hat{F}_{pts}^{t_0}$ serves as the key and value. Then we can denote our LiDAR-enhanced image feature generated by our uncertainty-aware Deformable transformer block as

$$\hat{F}_{img}^{enh} = \hat{F}_{img}^{t_0} + \text{DeformAttn}_{\{U_{pts}^{t_0}\}}(\hat{F}_{img}^{t_0}, \hat{F}_{pts}^{t_0}), \quad (6)$$

where the subscript $\{U_{pts}^{t_0}\}$ highlights the involvement of uncertainty values of the LiDAR feature in the deformable attention calculation. Specifically, after obtaining the base attention weights W (as in standard deformable attention, predicted from the query), the contribution of the LiDAR feature is further augmented by its uncertainty map $U_{pts}^{t_0}$:

$$\tilde{W} = W \cdot [1 - \text{softmax}(U_{pts}^{t_0})]. \quad (7)$$

This \tilde{W} then replaces the original attention weights W for attention calculation. Here, the softmax normalization is adopted to highlight the most unreliable components in the predicted feature. Through a similar process, we also obtain the image-enhanced feature \hat{F}_{pts}^{enh} .

During training, the extracted features $F_{img}^{t_0}$ and $F_{pts}^{t_0}$ from the backbone are used as supervision for the uncertainty-aware fusion, and we adopt an ℓ_2 loss to measure the discrepancy between the fusion features and the target:

$$\mathcal{L}_{\text{Fuse}} = \|\hat{F}_{img}^{enh} - F_{img}^{t_0}\|_2^2 + \|\hat{F}_{pts}^{enh} - F_{pts}^{t_0}\|_2^2. \quad (8)$$

C. Training and Inference Strategy.

During training, we optimize the two modules in separate stages. The motivation is to first establish reliable temporal dynamics before introducing cross-modality interaction, thereby preventing unstable temporal predictions from propagating severe noise into the fusion stage. We denote the detector's standard detection loss (comprising classification and regression terms) as \mathcal{L}_{det} . In the first stage, the *History-based Temporal Transformer* is trained jointly with $\mathcal{L}_{\text{TempPred}}$ and \mathcal{L}_{det} to capture temporal feature prediction:

$$\mathcal{L}_{\text{HFP}} = \mathcal{L}_{\text{TempPred}} + \mathcal{L}_{\text{det}}. \quad (9)$$

In the second stage, we freeze the temporal transformer to stabilize the learned temporal modeling, and then optimize the *Uncertainty-based Cross-modality Transformer* using $\mathcal{L}_{\text{Uncert}}$, $\mathcal{L}_{\text{Fuse}}$, and \mathcal{L}_{det} :

$$\mathcal{L}_{\text{UncertFuse}} = \mathcal{L}_{\text{Uncert}} + \mathcal{L}_{\text{Fuse}} + \mathcal{L}_{\text{det}}. \quad (10)$$

During inference, the history memory bank is continuously updated to maintain temporal continuity under arbitrary modality drop. If a modality is available, the extracted feature $F_M^{t_0-1}$ is stored in the bank; otherwise, the compensated feature $\hat{F}_M^{t_0-1}$ is used. This mechanism prevents disruption of the temporal stream when inputs are missing. For detection, the compensated feature $\hat{F}_M^{t_0}$ is used when a modality is absent, while in the normal case, the extracted feature $F_M^{t_0}$ is directly adopted. This strategy ensures robustness by preserving temporal consistency and fully exploiting reliable features under both complete and degraded sensor conditions.

TABLE I

PERFORMANCE COMPARISON IN mAP (%) AND NDS (%) OF DETECTORS AND THEIR MODALPATCH-ENHANCED COUNTERPARTS ON nuSCENES VALIDATION SET UNDER DROP_RATES = {10%, 30%, 50%}, WHERE “NO DROP” DENOTES THE SETTING WHERE ALL SENSOR MODALITIES ARE FULLY AVAILABLE AND VALUES IN PARENTHESES DENOTE THE PERFORMANCE BOOSTS OVER THE DETECTORS.

Methods	No drop		Drop_rate = 10%		Drop_rate = 30%		Drop_rate = 50%	
	mAP	NDS	mAP	NDS	mAP	NDS	mAP	NDS
UniBEV			59.66	65.47	48.86	58.53	35.49	50.65
+ModalPatch	64.24	68.53	62.65 (+2.99)	67.35 (+1.88)	56.56 (+7.70)	63.40 (+4.87)	46.32 (+10.83)	56.91 (+6.26)
BevFusion			60.74	67.46	43.26	58.45	27.74	50.37
+ModalPatch	68.73	71.47	63.37 (+2.63)	68.43 (+0.97)	50.59 (+7.33)	61.30 (+2.85)	31.43 (+3.69)	50.60 (+0.23)
CMT			60.51	67.81	42.90	58.47	27.21	49.92
+ModalPatch	70.07	72.68	67.96 (+7.45)	71.42 (+3.61)	58.77 (+15.87)	65.81 (+7.34)	44.21 (+17.00)	56.89 (+6.97)
MEFormer			62.93	69.40	45.37	60.29	27.88	50.64
+ModalPatch	71.49	73.86	68.65 (+5.72)	72.11 (+2.71)	59.02 (+13.65)	66.30 (+6.01)	44.11 (+16.23)	57.36 (+6.72)

IV. EXPERIMENT

A. Experiment Settings

Dataset and Metrics. We conduct experiments on the large-scale *nuScenes* dataset [9], which contains 1000 driving scenes (850 for training and 150 for validation) collected in Boston and Singapore. The dataset is captured using a 32-beam LiDAR and six surrounding cameras covering a full 360° field of view. We use the official training and validation splits for model training and evaluation, respectively. To simulate random modality-drop scenarios during testing, for each input modality $M \in \{\text{Camera, LiDAR}\}$ we apply a drop probability $\text{Drop_rate} \in \{0.1, 0.3, 0.5\}$ to independently discard the corresponding inputs. Since modality dropping is modeled as independent events, it is possible that both modalities are absent at the same time step. For model evaluation, we adopt the official metrics of *nuScenes* [9], the *mean Average Precision* (mAP) and *nuScenes Detection Score* (NDS), as primary evaluation metrics to assess detection performance.

Models and Implementation Details. Our proposed *ModalPatch* is, to the best of our knowledge, the first plug-and-play module for multi-modal 3D object detection under modality-drop scenarios, which can be seamlessly integrated into existing detectors. To demonstrate its generality, we evaluate *ModalPatch* on two representatives of detectors: the BEV-based BEVFusion [14], and the transformer-based CMT [10], together with their recent modality-missing extensions, UniBEV [11] and MEFormer [12]. In all experiments, we initialize the detectors with the officially released, pre-trained weights and keep these parameters frozen during the subsequent training. When applying *ModalPatch* in training, we remove the randomness in data augmentation from each detector to ensure feature consistency across consecutive frames. For training *ModalPatch*, we use AdamW as the optimizer with a learning rate of 0.0002. The training of both stages is conducted for 12 epochs. The temporal length of the memory bank is set to $\tau = 6$. During testing, for baseline detectors without *ModalPatch*, UniBEV includes a built-in mechanism to handle “None” inputs, and other detectors replace missing modalities with zero-filled tensors to ensure

inference execution. For detectors + *ModalPatch*, we follow the inference procedure described in Section III-C to handle modality drops. Both training and testing are conducted with a batch size of 1 on the NVIDIA RTX 3090 GPU.

B. Detection under Modality Drop

Table I demonstrates the superiority of *ModalPatch* in enhancing robustness under modality-drop scenarios. While all baseline detectors suffer severe degradation, *ModalPatch* consistently alleviates the performance collapse across diverse architectures. Under a moderate drop rate of 10%, *ModalPatch* provides consistent improvements across all detectors, with an average gain of +4.70% mAP and +2.29% NDS, effectively mitigating the performance degradation caused by missing modalities and keeping the performance closer to the no-missing upper bound. Under the 30% drop rate, our method yields substantial improvements (average +11.14% in mAP and +5.28% in NDS), showing the effectiveness of temporal feature prediction and uncertainty-guided fusion in recovering missing signals. Under the challenging 50% drop condition, *ModalPatch* still achieves significant gains, leading to an average improvement of +11.93% mAP and +5.05% NDS. Notably, UniBEV incorporates a feature-adaptive fusion strategy that provides relatively strong baseline performance under severe modality loss; however, *ModalPatch* still delivers a significant margin of improvement (e.g., +10.83% mAP / +6.26% NDS at 50% drop), demonstrating its complementary value. These results highlight that *ModalPatch* not only provides consistent benefits to both BEV-based and transformer-based detectors but also delivers the most significant advantages in extreme failure cases, confirming its practicality as a plug-and-play solution for real-world deployment. In cases of “no drop”, *ModalPatch* preserves the baseline performance by leaving original features unaltered.

Detection Visualization. Taking UniBEV and CMT for examples, Fig. 5 presents qualitative visualizations under the challenging 50% drop rate and details three cases: camera drop, LiDAR drop, and both-drop. In the camera drop case, the baselines fail to detect many distant ob-

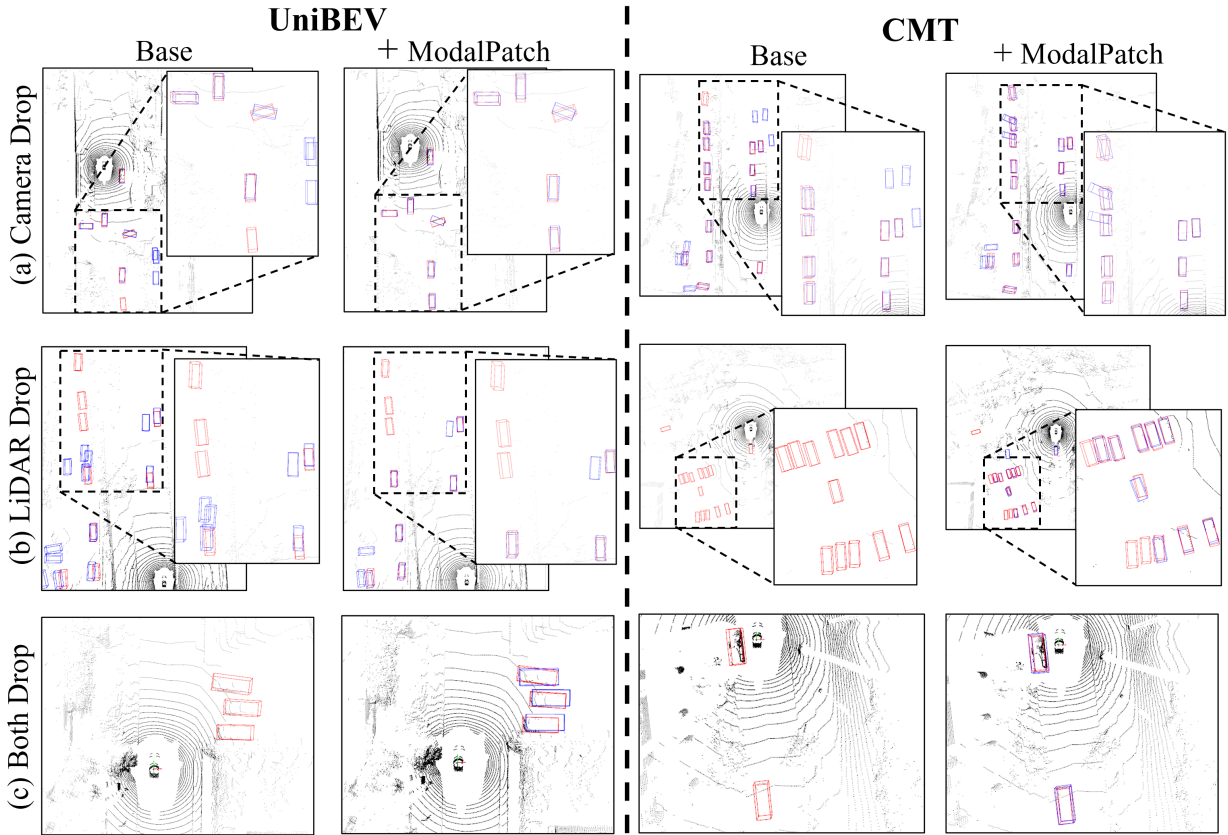


Fig. 5. Qualitative visualizations of UniBEV and CMT detectors w/ or w/o ModalPatch under a 50% drop rate. Each pair compares the baseline detector (Base) with the detector enhanced by ModalPatch (+ModalPatch), where red boxes denote ground-truth objects and blue ones denote detected objects.

jects because sparse LiDAR scans alone cannot provide sufficient appearance cues. With ModalPatch, the missing image information (e.g., color and texture) is effectively compensated, leading to improved recall of far-range objects and more precise bounding boxes. In the LiDAR drop case, UniBEV exhibits box localization biases due to the lack of depth information, while CMT—being heavily dependent on LiDAR—completely misses several targets. By contrast, with ModalPatch, the compensated LiDAR features restore spatial cues, resulting in more accurate localization and higher recall. Finally, in the extreme both-drop case, neither UniBEV nor CMT can produce meaningful detections because they lack mechanisms to handle simultaneous modality loss. With ModalPatch, however, the integration of temporal prediction and uncertainty-guided cross-modal fusion still enables the detectors to capture object instances, demonstrating strong robustness and reliability under extreme sensor failures.

C. Ablation Studies

We conduct ablation experiments primarily using UniBEV and CMT as the baseline detectors, under a 50% drop rate, to further analyze the proposed modules.

Component ablation. Table II analyses the individual contributions of our proposed modules. Starting from the baseline detector (a), adding the history-based feature prediction module (b) significantly boosts performance (e.g., +6.55% mAP / +3.72% NDS on UniBEV and +13.24% mAP / +4.93% NDS

TABLE II
COMPONENT ABLATION IN mAP/NDS (%) OF PROPOSED MODULES AT $DROP_RATES=50\%$, WHERE BASE DENOTES THE ORIGINAL DETECTOR.

	Base	HFP	UCF	UniBEV	CMT
(a)	✓			35.49/50.65	27.21/49.92
(b)	✓	✓		42.04/54.37	40.45/54.85
(c)	✓	✓	✓	46.32/56.91	44.21/56.89

on CMT), demonstrating that temporal feature modeling effectively compensates for missing modalities. Building on this, integrating the uncertainty-guided cross-modality fusion module (c) provides additional noticeable gains (e.g., +4.28% mAP / +2.54% NDS on UniBEV and +3.76% mAP / +2.04% NDS on CMT), as cross-modal interaction refines the compensated features and suppresses prediction bias. Together, the two modules work complementarily, achieving robust and reliable detection even under severe modality-missing conditions.

We also analyze the feature’s *mean squared error* (MSE) with respect to the ground-truth features (i.e., extracted features with no drop) for both camera and LiDAR modalities, as shown in Fig. 6. The top row presents the baseline results without ModalPatch, which exhibit large reconstruction errors due to the absence of features. With the history-based

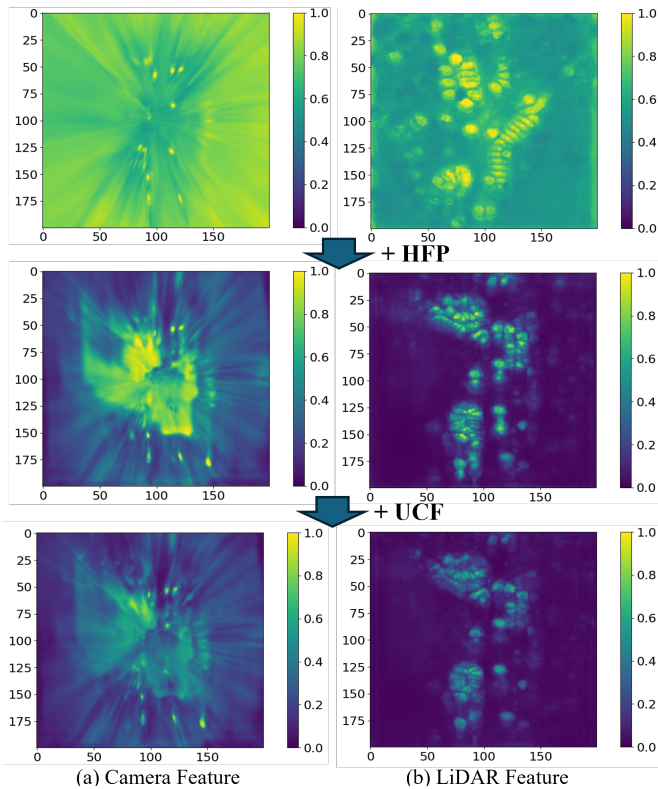


Fig. 6. Feature mean squared error between the predicted features and the ground-truth features for camera (a) and LiDAR (b) modalities using UniBEV with ModalPatch, with the top row for the baseline mechanism (i.e., zero-filled feature due to modality drop), the middle row for adopting the HFP, and the bottom row for continuously adopting the UCF.

feature prediction module (middle row), temporal modeling effectively reduces these errors by compensating for missing features. After further incorporating the uncertainty-guided cross-modality transformer (bottom row), the error maps are further suppressed and become more uniformly dim, indicating the effective suppression on the prediction bias. By complementarily adopting two proposed modules, the reliability of compensated features is enhanced, leading to more accurate feature recovery under missing modalities.

Ablation on uncertainty-guided cross-modality fusion. Table III compares the performance of the cross-modality transformer with and without the proposed uncertainty modeling. Without uncertainty guidance, the fusion relies purely on spatial attention, leading to moderate improvements but still propagating biased or unreliable features. By contrast, equipping the transformer with uncertainty estimation yields consistent gains (e.g., +1.39% mAP / +1.49% NDS on UniBEV and +1.02% mAP / +0.92% NDS on CMT), confirming that explicitly quantifying the reliability of compensated features allows the model to suppress biased signals and emphasize trustworthy information.

Ablation on single-modality conditions. Table IV evaluates the performance under the 50% drop rate when only a single modality (camera or LiDAR) is available. Since no complementary modality can be exploited in this setting, we only employ the HFP module to provide temporal

TABLE III
ABLATION OF THE PROPOSED UCF WITH AND WITHOUT UNCERTAINTY MODELING IN *mAP* (%) AND *NDS* (%) UNDER 50% MODALITY DROP.

Methods	UniBEV		CMT	
	mAP	NDS	mAP	NDS
w/o uncertainty modeling	44.93	55.42	43.19	55.97
w/ uncertainty modeling	46.32	56.91	44.21	56.89

TABLE IV
RESULTS IN *mAP* (%) AND *NDS* (%) UNDER THE 50% DROP RATE SCENARIO WHEN ONLY A SINGLE MODALITY (LiDAR-ONLY OR CAMERA-ONLY) IS AVAILABLE. SINCE NO COMPLEMENTARY MODALITIES CAN BE UTILIZED, ONLY THE HFP MODULE IS ADOPTED FOR TEMPORAL COMPENSATION.

Methods	LiDAR-only		Camera-only	
	mAP	NDS	mAP	NDS
UniBEV	23.85	47.58	13.45	31.07
+HFP	36.62	52.50	24.77	36.48
No modality drop	58.16	65.26	35.00	42.40
CMT	24.13	48.37	0.02	1.69
+HFP	27.82	49.78	0.17	5.03
No modality drop	60.01	67.23	0.26	8.84

compensation. The results show that HFP brings substantial improvements for UniBEV, with gains of +12.77% mAP / +4.92% NDS in the LiDAR-only case and +11.32% mAP / +5.41% NDS in the camera-only case, demonstrating its strong effectiveness in alleviating single-modality collapse. For CMT, HFP also improves the LiDAR-only performance by +3.69% mAP / +1.41% NDS. In the camera-only case, the gain is more modest, as the detector shows a strong reliance on LiDAR and its performance ceiling with no camera input missing is inherently low (0.26% mAP / 8.84% NDS). These results demonstrate that our method continues to yield measurable improvements, suggesting that temporal modeling can enhance robustness and preserve detection capability even with only a single modality existing.

Runtime Speed. Table V shows the runtime speed measured in frames per second (FPS) with a 50% drop rate. On average, integrating ModalPatch results in an average runtime reduction from 5.33 FPS to 4.90 FPS, while providing substantial robustness and accuracy improvements under modality-missing conditions as in Table I. This demonstrates that ModalPatch achieves a favorable trade-off between detection performance and computational efficiency, making it practical for real-world deployments with modality drop.

D. Limitation

Despite the effectiveness of ModalPatch in improving robustness under modality-drop scenarios, its benefits are constrained when the detector performs poorly on a single modality with inherently low detection capability, as our

TABLE V

RUNTIME SPEED MEASURED WITH FRAMES PER SECOND (FPS) UNDER
DROP_RATE=50% ON NVIDIA RTX 3090 GPU.

Methods	UniBEV	BEVFusion	CMT	MEFormer	Ave
Base	3.21	8.03	4.53	5.56	5.33
+ModalPatch	3.06	7.54	4.37	4.62	4.90

method primarily compensates for the unfavorable modality's features through history-based feature prediction. In the future, we plan to explore stronger single-modality feature enhancement to further overcome this limitation.

V. CONCLUSION

In this paper, we present ModalPatch, a plug-and-play robustness module for multi-modal 3D object detection, which aims to enable existing detectors to cope with independent sensor drops without requiring significant retraining or architectural modifications. ModalPatch consists of two complementary components: a history-based feature prediction module that leverages temporal continuity to model the evolution of features across consecutive frames and predict current representations for missing inputs; and an uncertainty-guided cross-modality fusion module that estimates the reliability of the compensated features and performs spatially-aware cross-modal fusion, suppressing biased signals while reinforcing trustworthy information. Extensive experiments demonstrate that ModalPatch consistently improves the robustness of diverse detectors under various modality-drop conditions.

ACKNOWLEDGMENT

This work was supported in part by Canada CIFAR AI Chairs Program, the Natural Sciences and Engineering Research Council of Canada, JST CRONOS Grant (No. JPMJCS24K8), JSPS KAKENHI Grant (No.JP21H04877, No.JP23H03372, and No.JP24K02920).

REFERENCES

- [1] Y. Yuan, Z. Li, and B. Zhao, "A survey of multimodal learning: Methods, applications, and future," *ACM Computing Surveys*, vol. 57, no. 7, pp. 1–34, 2025.
- [2] X. Zhang, H. Wang, and H. Dong, "A survey of deep learning-driven 3d object detection: Sensor modalities, technical architectures, and applications," *Sensors*, vol. 25, no. 12, p. 3668, 2025.
- [3] C. Liao, K. Lei, X. Zheng, J. Moon, Z. Wang, Y. Wang, D. P. Paudel, L. Van Gool, and X. Hu, "Benchmarking multi-modal semantic segmentation under sensor failures: Missing and noisy modality robustness," in *Proceedings of the Computer Vision and Pattern Recognition Conference*, 2025, pp. 1576–1586.
- [4] M. Bijelic, T. Gruber, and W. Ritter, "A benchmark for lidar sensors in fog: Is detection breaking down?" in *2018 IEEE intelligent vehicles symposium (IV)*. IEEE, 2018, pp. 760–767.
- [5] F. Chang, E. Jafarzadeh, J. Del Gatto, G. Cran, and H. Sadjadi, "Failure mode investigation to enable lidar health monitoring for automotive application," in *Annual Conference of the PHM Society*, vol. 15, no. 1, 2023.
- [6] T. Goelles, B. Schlager, and S. Muckenhuber, "Fault detection, isolation, identification and recovery (fdiir) methods for automotive perception sensors including a detailed literature survey for lidar," *Sensors*, vol. 20, no. 13, p. 3662, 2020.
- [7] M. Bijelic, F. Mannan, T. Gruber, W. Ritter, K. Dietmayer, and F. Heide, "Seeing through fog without seeing fog: Deep sensor fusion in the absence of labeled training data," *arXiv preprint arXiv:1902.08913*, vol. 1030, 2019.
- [8] Q. Xu, Y. Zhong, and U. Neumann, "Behind the curtain: Learning occluded shapes for 3d object detection," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 3, 2022, pp. 2893–2901.
- [9] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusenes: A multimodal dataset for autonomous driving," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 11 621–11 631.
- [10] J. Yan, Y. Liu, J. Sun, F. Jia, S. Li, T. Wang, and X. Zhang, "Cross modal transformer: Towards fast and robust 3d object detection," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 18 268–18 278.
- [11] S. Wang, H. Caesar, L. Nan, and J. F. Kooij, "Unibev: Multi-modal 3d object detection with uniform bev encoders for robustness against missing sensor modalities," in *2024 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2024, pp. 2776–2783.
- [12] J. Cha, M. Joo, J. Park, S. Lee, I. Kim, and H. J. Kim, "Robust multimodal 3d object detection via modality-agnostic decoding and proximity-based modality ensemble," *arXiv preprint arXiv:2407.19156*, 2024.
- [13] Z. Song, L. Liu, F. Jia, Y. Luo, C. Jia, G. Zhang, L. Yang, and L. Wang, "Robustness-aware 3d object detection in autonomous driving: A review and outlook," *IEEE Transactions on Intelligent Transportation Systems*, 2024.
- [14] Z. Liu, H. Tang, A. Amini, X. Yang, H. Mao, D. Rus, and S. Han, "Bevfusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," *arXiv preprint arXiv:2205.13542*, 2022.
- [15] T. Liang, H. Xie, K. Yu, Z. Xia, Z. Lin, Y. Wang, T. Tang, B. Wang, and Z. Tang, "Bevfusion: A simple and robust lidar-camera fusion framework," *Advances in Neural Information Processing Systems*, vol. 35, pp. 10 421–10 434, 2022.
- [16] H. Cai, Z. Zhang, Z. Zhou, Z. Li, W. Ding, and J. Zhao, "Bevfusion4d: Learning lidar-camera fusion under bird's-eye-view via cross-modality guidance and temporal aggregation," *arXiv preprint arXiv:2303.17099*, 2023.
- [17] Y. Liu, T. Wang, X. Zhang, and J. Sun, "Petr: Position embedding transformation for multi-view 3d object detection," in *European conference on computer vision*. Springer, 2022, pp. 531–548.
- [18] X. Bai, Z. Hu, X. Zhu, Q. Huang, Y. Chen, H. Fu, and C.-L. Tai, "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.
- [19] H. Zhang, L. Liang, P. Zeng, X. Song, and Z. Wang, "Sparselif: High-performance sparse lidar-camera fusion for 3d object detection," in *European Conference on Computer Vision*. Springer, 2024, pp. 109–128.
- [20] B. Sinopoli, L. Schenato, M. Franceschetti, K. Poolla, M. I. Jordan, and S. S. Sastry, "Kalman filtering with intermittent observations," *IEEE transactions on Automatic Control*, vol. 49, no. 9, pp. 1453–1464, 2004.
- [21] Z. Zong, D. Jiang, G. Song, Z. Xue, J. Su, H. Li, and Y. Liu, "Temporal enhanced training of multi-view 3d object detector via historical object prediction," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 3781–3790.
- [22] Z. Xia, X. Pan, S. Song, L. E. Li, and G. Huang, "Vision transformer with deformable attention," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 4794–4803.
- [23] C.-P. Lee and C.-J. Lin, "A study on l2-loss (squared hinge-loss) multiclass svm," *Neural computation*, vol. 25, no. 5, pp. 1302–1323, 2013.
- [24] A. Kendall and Y. Gal, "What uncertainties do we need in bayesian deep learning for computer vision?" *Advances in neural information processing systems*, vol. 30, 2017.
- [25] K. Wang, C. Shen, X. Li, and J. Lu, "Uncertainty quantification for safe and reliable autonomous vehicles: A review of methods and applications," *IEEE Transactions on Intelligent Transportation Systems*, 2025.
- [26] A. Kumar, T. K. Marks, W. Mou, C. Feng, and X. Liu, "Uglli face alignment: Estimating uncertainty with gaussian log-likelihood loss," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 0–0.