

Multimodal Mixture-of-Experts for ISAC in Low-Altitude Wireless Networks

Kai Zhang, *Graduate Student Member, IEEE*, Wentao Yu, *Member, IEEE*, Hengtao He, *Member, IEEE*, Shenghui Song, *Senior Member, IEEE*, Jun Zhang, *Fellow, IEEE*, and Khaled B. Letaief, *Fellow, IEEE*

Abstract—Integrated sensing and communication (ISAC) is a key enabler for low-altitude wireless networks (LAWNs), providing simultaneous environmental perception and data transmission in complex aerial scenarios. By combining heterogeneous sensing modalities such as visual, radar, lidar, and positional information, multimodal ISAC can improve both situational awareness and robustness of LAWNs. However, most existing multimodal fusion approaches use static fusion strategies that treat all modalities equally and cannot adapt to channel heterogeneity or time-varying modality reliability in dynamic low-altitude environments. To address this fundamental limitation, we propose a mixture-of-experts (MoE) framework for multimodal ISAC in LAWNs. Each modality is processed by a dedicated expert network, and a lightweight gating module adaptively assigns fusion weights according to the instantaneous informativeness and reliability of each modality. To improve scalability under the stringent energy constraints of aerial platforms, we further develop a sparse MoE variant that selectively activates only a subset of experts, thereby reducing computation overhead while preserving the benefits of adaptive fusion. Comprehensive simulations on three typical ISAC tasks in LAWNs demonstrate that the proposed frameworks consistently outperform conventional multimodal fusion baselines in terms of learning performance and training sample efficiency.

Index Terms—6G, integrated sensing and communication (ISAC), low-altitude wireless networks (LAWNs), mixture-of-experts (MoE), multimodal learning.

I. INTRODUCTION

The rapid advancement of sixth-generation (6G) wireless networks is poised to revolutionize communication systems, shifting from conventional two-dimensional terrestrial networks to three-dimensional ubiquitous connectivity [2], [3]. Within this landscape, low-altitude wireless networks (LAWNs) are emerging as a critical infrastructure. By deploying unmanned aerial vehicles (UAVs) as flexible wireless relays, LAWNs form resilient three-dimensional communication links that bridge the coverage gaps of fixed terrestrial

infrastructure [4], [5]. This architecture enables rapid deployment of wireless coverage in disaster-affected and remote areas where fixed infrastructure is unavailable or damaged [6]. Consequently, LAWNs are increasingly used to support applications such as intelligent aerial logistics [7], precision agriculture [8], and large-scale infrastructure monitoring [9], [10].

Despite great potentials, LAWNs must maintain reliable wireless links while simultaneously performing real-time environmental perception in congested low-altitude airspace. To tackle this challenge, integrated sensing and communication (ISAC) has emerged as a new paradigm that promotes mutual benefits between wireless communication and sensing within a unified system [11], [12]. This is particularly important for LAWNs, where the complex aerial conditions render isolated communication or sensing unreliable. Specifically, the sensed spatial information can stabilize wireless links in high-mobility scenarios by enabling proactive beam alignment and blockage prediction. In addition, reliable connectivity enables real-time fusion of distributed sensing data, extending the sensing coverage of individual UAVs.

Given these promising capabilities, extensive research has investigated the integration of ISAC into LAWN systems, including advanced beamforming techniques [13], UAV trajectory optimization [14], and spectrum-efficient resource allocation [15]. Specifically, a substantial body of works have investigated sensing-aided communication, utilizing environmental awareness to facilitate robust UAV connectivity in high-mobility aerial scenarios [16]–[19]. For instance, by leveraging historical sensing data and extended Kalman filtering (EKF) to predict target motion, Zhou *et al.* [16] enabled precise beamforming without pilots to boost communication throughput, while Jiang *et al.* [17] optimized UAV trajectory to guarantee link maintenance. Focusing on reliability, Jiang *et al.* [18] maximized communication outage capacity by optimizing the sensing duration and UAV trajectory based on the outage probabilities predicted from sensing data. Furthermore, Meng *et al.* [19] proposed a two-stage sensing protocol to acquire target location information, thereby optimizing the phase shifts of intelligent surface to maximize the achievable communication rate. Conversely, other studies explore communication-aided sensing, exploiting ubiquitous wireless infrastructure to extend environmental perception coverage and precision [20]–[22]. For instance, Ji *et al.* [20] utilized existing long-term evolution (LTE) communication signals to track drone trajectories by analyzing Doppler frequency shifts, demonstrating that communication waveforms can be repurposed for high-

This work was presented in part at IEEE International Symposium on Personal, Indoor, Mobile Radio Communications (PIMRC), Istanbul, Turkey, Sept. 2025 [1].

Kai Zhang, Shenghui Song, Jun Zhang, and Khaled B. Letaief are with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong (email: kzhangbn@connect.ust.hk, eeshsong@ust.hk, eejzhang@ust.hk, eekhaled@ust.hk).

Wentao Yu was with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong. He is now with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: wentaoyu@ece.ubc.ca).

Hengtao He is with the National Mobile Communications Research Laboratory, Southeast University, Nanjing, China (hehengtao@seu.edu.cn).

precision sensing without dedicated hardware. Extending this to networked cooperation, Yang *et al.* [21] proposed fusing signals from distributed communication nodes to improve sensing accuracy. To further exploit the potential of infrastructure cooperation, Lu *et al.* [22] developed a symbol-level fusion method that utilizes the distributed network to achieve superior UAV localization and velocity estimation compared to single-station approaches. While extensive studies have advanced the development of ISAC in LAWN systems, they primarily focused on unimodal sensing that leverages only radio frequency (RF) signals, extracting environmental features solely from communication signals such as channel state information (CSI), received signal strength, or Doppler shifts. Although RF-based sensing is ubiquitous, it provides only coarse environmental perception and is sensitive to blockage, multipath scattering, and low signal-to-noise ratios (SNRs). Consequently, reliance on a single sensing modality leads to performance degradation under high-speed aerial dynamics, compromising both the wireless connectivity and perception reliability essential for LAWNs.

To mitigate these limitations, multimodal ISAC integrates complementary sensing modalities to enhance both environmental perception and communication reliability in LAWNs [23], [24]. By synergizing heterogeneous sensing sources (e.g., RF signals, visual images, and lidar point clouds), this approach exploits the unique strengths of different modalities to construct a high-fidelity spatial representation. For instance, the weather resilience of radar can effectively offset visual degradation in adverse conditions, while the high angular resolution of lidar resolves the spatial ambiguity typical of RF signals. Such fusion improves situational awareness and robustness to modality-specific failures that occur in dynamic aerial scenarios. Several recent works have investigated the design of multimodal ISAC frameworks, developing different fusion architectures to align and integrate heterogeneous sensing modalities [25]–[29]. For instance, Zhu *et al.* [25] proposed a cross-modal feature enhancement mechanism that fuses image and radar data for multimodal beam prediction tasks, demonstrating improved robustness under varying modality quality. Moving towards more unified representations, Tariq *et al.* [26] and Cui *et al.* [27] explored transformer-based architectures that fuse multiple sensing modalities for accurate beam prediction and reliable communication in ISAC systems. To further address the generalization limitations of task-specific models, Cheng *et al.* [28] introduced multimodal large language models (MLLMs) into ISAC, leveraging their scalable cross-modal capabilities. Specific to the low-altitude airspace, Xie *et al.* [29] developed a transformer-based perception fusion network that aligns vision, lidar, and radar data for proactive beamforming in dynamic LAWN scenarios.

However, most existing multimodal fusion methods employ static fusion strategies that treat all modalities equally and cannot adapt to rapid environmental changes or time-varying modality reliability in low-altitude environments. To address this limitation, we propose a multimodal mixture-of-experts (MoE) architecture that enables adaptive cross-modal fusion for ISAC in LAWNs. Unlike static baselines, the proposed architecture orchestrates a set of modality-specific expert

networks through an adaptive gating mechanism that assigns data-dependent fusion weights to each modality. This design allows the system to assess the instantaneous informativeness of each modality and place greater weight on reliable signals while down-weighting degraded ones. As a result, the fused representation becomes more robust to large fluctuations in sensing quality on high-speed aerial platforms in LAWNs.

A. Contributions

The main contributions of this paper are summarized as follows.

1. We establish a unified multimodal learning framework for ISAC in LAWNs by modeling the relationship between heterogeneous sensing modalities and RF communication signals as a generalized cross-modal prediction problem. This formulation provides a common optimization paradigm that seamlessly supports diverse perception and communication tasks in low-altitude ISAC systems, including sensing-aided beam prediction, channel estimation, and communication-aided UAV trajectory tracking.

2. We propose a novel multimodal MoE architecture that orchestrates modality-specific expert networks through an adaptive gating mechanism, enabling dynamic context-aware fusion based on the instantaneous informativeness and reliability of each modality. To meet the energy and computation constraints of aerial platforms, we further develop a sparse MoE variant that activates only a subset of experts during inference while preserving the benefits of adaptive fusion.

3. We conduct extensive simulations on representative ISAC tasks in realistic LAWN scenarios, including sensing-aided beam prediction, sensing-aided path loss prediction, and communication-aided UAV trajectory tracking. Simulation results demonstrate that the proposed multimodal MoE and sparse MoE frameworks consistently outperform conventional multimodal fusion baselines in terms of learning performance and training sample efficiency.

B. Organization and Notations

The remainder of this paper is organized as follows. Section II elaborates on the system model and formulates the unified learning framework for ISAC in LAWNs. Section III presents the proposed multimodal MoE architecture. Section IV introduces the sparse MoE variant to address the resource constraints of aerial platforms. Simulation results are presented in Section V, and we conclude the paper in Section VI.

Notations: Column vectors and matrices are denoted by boldface lowercase and boldface capital letters, respectively. The symbol $\mathbb{R}$ denotes the set of real numbers. $\mathbb{C}^{M \times N}$ represents the space of the $M \times N$ complex-valued matrices. $(\cdot)^T$ and $(\cdot)^H$ stand for the transpose and the conjugate transpose of their arguments, respectively. $\text{Tr}(\mathbf{A})$ denotes the trace of matrix $\mathbf{A}$. $\mathbb{E}[\cdot]$ denotes the expectation operation. ∇ represents the gradient operator. $|\cdot|$ and $\|\cdot\|$ stand for the ℓ_1 and ℓ_2 norm of vectors, respectively.

II. SYSTEM MODEL AND PROBLEM FORMULATION

In this section, we first elaborate on the multimodal ISAC system model in LAWNs. We then establish a unified multimodal learning framework that incorporates heterogeneous




Fig. 1. An illustration of multimodal ISAC in LAWNs, where a ground BS and UAVs exploit heterogeneous sensing data to support sensing-aided beam prediction, sensing-aided channel estimation, and communication-aided UAV trajectory tracking.

sensing modalities and RF communication signals to support diverse perception and communication tasks. Finally, we instantiate this framework through three representative ISAC tasks to demonstrate its broad applicability.

A. Multimodal ISAC System Model

We consider a multimodal ISAC system deployed in LAWNs, as shown in Fig. 1. The system comprises a ground base station (BS) equipped with M antennas and K single-antenna authorized UAVs, denoted by the set $\mathcal{K} = \{1, \dots, K\}$. The BS conducts downlink communication with the UAVs while simultaneously acquiring spatial and geometric information of the propagation environment. In addition to RF transceivers, both the BS and UAVs are equipped with heterogeneous sensors, such as cameras, lidars, radars, and global positioning systems (GPS). These heterogeneous sensing platforms collectively provide multimodal observations of the surrounding environment, which are exploited to enhance both environmental perception and communication performance. The communication and sensing models are detailed as follows.

1) *Communication Model*: We consider a frequency-division duplexing millimeter-wave (mmWave) communication system operating over a flat block-fading channel. For the downlink data transmission, let $\mathbf{V}(t) = [\mathbf{v}_1(t), \dots, \mathbf{v}_K(t)] \in \mathbb{C}^{M \times K}$ denote the beamforming matrix at the BS in time slot t , where each column vector $\mathbf{v}_k(t) \in \mathbb{C}^{M \times 1}$ corresponds to the beamforming vector designated for the k -th UAV. Then, the transmitted signal vector is given by

$$\mathbf{x}(t) = \sum_{k=1}^K \mathbf{v}_k(t) s_k(t) = \mathbf{V}(t) \mathbf{s}(t), \quad (1)$$

where $\mathbf{s}(t) = [s_1(t), \dots, s_K(t)]^T \in \mathbb{C}^{K \times 1}$ is the vector of transmitted symbols with $\mathbb{E}[|s_k(t)|^2] = 1$. The BS transmission power is constrained by

$$\text{Tr}(\mathbf{V}(t) \mathbf{V}^H(t)) \leq P_{\text{DL}}, \quad (2)$$

where P_{DL} denotes the maximum downlink transmission power.

Let $\mathbf{h}_k(t) \in \mathbb{C}^{M \times 1}$ denote the downlink channel between the BS and UAV k at time slot t . We adopt a geometric multipath channel model [30], under which the downlink channel is expressed as

$$\mathbf{h}_k(t) = \sum_{\ell=1}^{L_p} \alpha_{\ell,k}(t) e^{-j2\pi f_c \tau_{\ell,k}(t)} \mathbf{a}(\theta_{\ell,k}(t)), \quad (3)$$

where L_p is the number of propagation paths, f_c is the carrier frequency, $\alpha_{\ell,k}(t)$ is the path loss of the ℓ -th path, $\tau_{\ell,k}(t) = d_{\ell,k}(t)/c$ is the associated propagation delay with $d_{\ell,k}(t)$ denoting the path length and c the speed of light, and $\theta_{\ell,k}(t)$ is the angle of departure (AoD). For a uniform linear array (ULA) with M antenna elements and inter-element spacing d_{ant} , the array response vector corresponding to AoD $\theta_{\ell,k}(t)$ is given by

$$\mathbf{a}(\theta_{\ell,k}(t)) = [1, e^{j\frac{2\pi}{\lambda} d_{\text{ant}} \sin(\theta_{\ell,k}(t))}, \dots, e^{j\frac{2\pi}{\lambda} d_{\text{ant}} (M-1) \sin(\theta_{\ell,k}(t))}]^T, \quad (4)$$

where λ is the carrier wavelength. Accordingly, the received signal at UAV k at time slot t can be written as

$$y_k(t) = \mathbf{h}_k^H(t) \mathbf{v}_k(t) s_k(t) + \sum_{j \neq k} \mathbf{h}_k^H(t) \mathbf{v}_j(t) s_j(t) + z_k(t), \quad (5)$$

where $z_k(t) \sim \mathcal{CN}(0, \sigma_0^2)$ denotes the additive white Gaussian noise with variance σ_0^2 . Consequently, the instantaneous downlink sum-rate is expressed as

$$R(t) = \sum_{k=1}^K \log_2 \left(1 + \frac{|\mathbf{h}_k^H(t) \mathbf{v}_k(t)|^2}{\sum_{j \neq k} |\mathbf{h}_k^H(t) \mathbf{v}_j(t)|^2 + \sigma_0^2} \right). \quad (6)$$

2) *Multimodal Sensing Model*: To facilitate robust communication and precise environmental perception within dynamic low-altitude scenarios, the multimodal ISAC framework synergizes heterogeneous sensing modalities, capturing complementary environmental features that augment the limited spatial awareness provided by standard RF communication

signals. Formally, at time slot t , the system aggregates a set of heterogeneous sensing inputs, denoted as

$$\mathbf{S}(t) = \{\mathbf{S}_{\text{vis}}(t), \mathbf{S}_{\text{radar}}(t), \mathbf{S}_{\text{lidar}}(t), \mathbf{S}_{\text{pos}}(t)\}, \quad (7)$$

where $\mathbf{S}_{\text{vis}}(t)$, $\mathbf{S}_{\text{radar}}(t)$, $\mathbf{S}_{\text{lidar}}(t)$, and $\mathbf{S}_{\text{pos}}(t)$ denote the sensing data from the vision, radar, lidar, and position modalities, respectively. These data are obtained from sensors either mounted on the UAVs or co-located at the ground BS. While $\mathbf{S}(t)$ is defined as a superset of potential sensing sources, our framework is inherently flexible and can operate with available subset of these modalities, depending on specific hardware configurations or instantaneous sensor availability. The specific characteristics and data representations of each modality are detailed as follows.

- Visual modality $\mathbf{S}_{\text{vis}}(t)$: The vision modality leverages RGB images to capture rich semantic and contextual information, such as the spatial distribution of obstacles and dynamic targets. Such visual data provide critical environmental awareness for downstream ISAC tasks like beam prediction and trajectory tracking. To optimize learning efficiency, the raw visual data is preprocessed through spatial filtering and background subtraction to suppress static environmental noise and isolate dynamic regions of interest (e.g., moving UAVs), thereby distilling the most task-relevant features for subsequent inference.
- Radar modality $\mathbf{S}_{\text{radar}}(t)$: The radar modality leverages Doppler signatures to provide precise measurements of target distance and relative velocity. Its primary advantage lies in the inherent robustness against adverse environmental conditions, such as fog, rain, or low light conditions, where optical sensors may fail. By encapsulating these kinematic characteristics, this modality offers critical motion information that enables the ISAC system to accurately track dynamic targets and anticipate mobility patterns even in visually degraded scenarios.
- Lidar modality $\mathbf{S}_{\text{lidar}}(t)$: Providing high-resolution three-dimensional geometric information, the lidar modality captures precise depth perception and spatial structural details that two-dimensional visual sensors inherently lack. By offering a fine-grained volumetric representation of the environment, this modality facilitates accurate spatial reasoning, enabling the system to resolve complex object shapes and determine precise physical locations essential for environment-aware communication and sensing tasks.
- Position modality $\mathbf{S}_{\text{pos}}(t)$: Providing a macroscopic spatial reference, the position modality supplies absolute localization coordinates and long-term motion trends. Unlike local environmental sensors, the position data offers a global context for user movement, serving as a critical coarse-grained information to anchor spatial relationships and guide decision-making in large-scale dynamic scenarios.

Collectively, the integration of these sensing modalities enables the construction of a comprehensive and task-relevant environmental representation. This multi-dimensional awareness is pivotal for supporting the dual objectives of robust communication and precise sensing in dynamic low-altitude environments.

B. Problem Formulation

To achieve robust perception and connectivity in dynamic LAWNs, we formulate a unified multimodal learning framework that models the interaction between heterogeneous sensing modalities and RF communication signals as a generalized cross-modal prediction problem. This formulation unifies diverse ISAC tasks under a common optimization paradigm, enabling mutual performance enhancement by utilizing environmental perception to optimize connectivity or leveraging wireless signals to improve sensing precision.

Formally, at each time slot t , the multimodal ISAC system aggregates heterogeneous sensing inputs $\mathbf{S}(t)$ from different sensing modalities such as vision, radar, and lidar. In parallel, it incorporates RF link measurements $\mathbf{C}(t)$, which encompass wireless metrics including historical beam indices, received signal strength indicators, and CSI. To accommodate diverse task requirements, we define a unified system input $\mathbf{X}(t)$ by selectively aggregating heterogeneous sensing inputs $\mathbf{S}(t)$ and RF link measurements $\mathbf{C}(t)$. The resulting input space is defined as

$$\mathbf{X}(t) \in \{\mathbf{S}(t), \mathbf{C}(t), \{\mathbf{S}(t), \mathbf{C}(t)\}\}, \quad (8)$$

which allows the predictive model to operate on either individual domains or a composite measurement. We aim to learn a parametric mapping function g_{Θ} that maps the system input $\mathbf{X}(t)$ to a task-specific output $\mathbf{o}(t)$, such as a discrete beam index or a continuous UAV coordinate. The learning objective is formulated as an expected risk minimization problem, where we seek the optimal parameters Θ that minimize the expected loss $\mathcal{L}(\hat{\mathbf{o}}(t), \mathbf{o}(t))$ between the model output $\hat{\mathbf{o}}(t)$ and the corresponding ground truth $\mathbf{o}(t)$ as follows,

$$\begin{aligned} \min_{\Theta} \quad & \mathbb{E}[\mathcal{L}(\hat{\mathbf{o}}(t), \mathbf{o}(t))], \\ \text{s.t.} \quad & \hat{\mathbf{o}}(t) = g_{\Theta}(\mathbf{X}(t)), \end{aligned} \quad (9)$$

where the expectation is taken over the spatiotemporal variability of the environment, channel states, and sensor observations.

This generalized formulation provides a versatile foundation applicable to a broad spectrum of ISAC scenarios. By customizing the output space and loss function, the framework seamlessly applies to diverse problem paradigms. For instance, defining $\mathbf{o}(t)$ as a beam index under cross-entropy loss recovers the beam prediction formulation, while setting $\mathbf{o}(t)$ as a position vector under mean squared error (MSE) yields the UAV trajectory tracking problem. To demonstrate this adaptability, we formally instantiate this framework through three representative ISAC tasks.

1) *Sensing-Aided Channel Estimation*: Precise acquisition of CSI is pivotal for mmWave communications but poses significant challenges in LAWNs due to the channel variations caused by rapid mobility. To mitigate the overhead of dense pilot signaling, we leverage heterogeneous multimodal sensing data $\mathbf{S}(t)$ as auxiliary side information to estimate the downlink channel vector $\mathbf{h}_k(t)$. To specialize the general Problem (9) for this task, we instantiate the system input $\mathbf{X}(t)$ as the multimodal sensing data $\mathbf{S}(t)$ and the model output $\hat{\mathbf{o}}(t)$ as the estimated downlink channel vector $\hat{\mathbf{h}}_k(t)$, i.e.,

$$\mathbf{X}(t) \triangleq \mathbf{S}(t) \quad \text{and} \quad \hat{\mathbf{o}}(t) \triangleq \hat{\mathbf{h}}_k(t). \quad (10)$$

Consequently, the optimization objective is concretized as a regression problem that minimizes the expected MSE between the estimated channel $\hat{\mathbf{h}}_k(t)$ and the ground truth $\mathbf{h}_k(t)$, given by

$$\begin{aligned} \min_{\Theta} \quad & \mathbb{E} \left[\left\| \mathbf{h}_k(t) - \hat{\mathbf{h}}_k(t) \right\|^2 \right] \\ \text{s.t.} \quad & \hat{\mathbf{h}}_k(t) = g_{\Theta}(\mathbf{S}(t)). \end{aligned} \quad (11)$$

By exploiting the rich spatial and semantic context provided by UAV-mounted and BS-mounted sensors, this formulation effectively captures the implicit correlations between environmental features and signal propagation, thereby enhancing CSI estimation accuracy while reducing the dependence on conventional pilot-based methods.

2) *Sensing-Aided Beamforming Design*: Efficient beam management is paramount for UAV-enabled mmWave communication systems, yet traditional beam sweeping incurs prohibitive latency overheads given the highly directional transmission and rapid aerial mobility. To enable low-latency beam management, we instantiate the unified Problem (9) as a proactive beam prediction task, which leverages heterogeneous sensing data to directly predict the optimal beamforming configuration. Specifically, the system input $\mathbf{X}(t)$ and output $\hat{\mathbf{o}}(t)$ are respectively instantiated as the multimodal sensing data $\mathbf{S}(t)$ and the estimated beamforming vectors $\{\hat{\mathbf{v}}_k(t)\}_{k=1}^K$ for all UAVs, i.e.,

$$\mathbf{X}(t) \triangleq \mathbf{S}(t) \quad \text{and} \quad \hat{\mathbf{o}}(t) \triangleq \{\hat{\mathbf{v}}_k(t)\}_{k=1}^K. \quad (12)$$

The general learning objective is specialized to maximize the instantaneous downlink sum-rate as follows,

$$\begin{aligned} \max_{\Theta} \quad & \mathbb{E} \left[\sum_{k=1}^K \log_2 \left(1 + \frac{|\mathbf{h}_k^H(t) \hat{\mathbf{v}}_k(t)|^2}{\sum_{j \neq k} |\mathbf{h}_k^H(t) \hat{\mathbf{v}}_j(t)|^2 + \sigma_0^2} \right) \right] \\ \text{s.t.} \quad & \hat{\mathbf{v}}_k(t) = g_{\Theta}(\mathbf{S}(t)), \forall k, \end{aligned} \quad (13)$$

where $\mathbf{h}_k(t)$ denotes the downlink channel vector between the BS and UAV k , and σ_0^2 is the noise power. By bypassing explicit channel estimation and exhaustive search, this approach effectively leverages the spatial intelligence embedded in multimodal observations to achieve accurate, mobility-aware beam selection, thereby stabilizing connectivity against the dynamic fluctuations of the low-altitude channel.

3) *Communication-Aided UAV Trajectory Tracking*: Accurate trajectory tracking is indispensable for maintaining reliable connectivity in LAWNs, yet relying solely on onboard sensors often leads to the degradation of localization accuracy due to signal blockage or inertial drift. To address this challenge, we employ onboard sensing data $\mathbf{S}(t)$ with downlink RF measurements $\mathbf{C}(t)$ to achieve high-precision positioning. Specifically, the system input $\mathbf{X}(t)$ and output $\hat{\mathbf{o}}(t)$ in Problem (9) are respectively instantiated as the joint sensing and RF measurements $\{\mathbf{S}(t), \mathbf{C}(t)\}$ and estimated UAV position $\hat{\mathbf{p}}(t)$, i.e.,

$$\mathbf{X}(t) \triangleq \{\mathbf{S}(t), \mathbf{C}(t)\} \quad \text{and} \quad \hat{\mathbf{o}}(t) \triangleq \hat{\mathbf{p}}(t). \quad (14)$$

The optimization objective is concretized as a trajectory regression problem that minimizes the cumulative position error

over the time horizon as follows,

$$\begin{aligned} \min_{\Theta} \quad & \mathbb{E} \left[\sum_{\tau=t-L}^t \|\mathbf{p}(\tau) - \hat{\mathbf{p}}(\tau)\|^2 \right] \\ \text{s.t.} \quad & \hat{\mathbf{p}}(\tau) = g_{\Theta}(\{\mathbf{S}(\tau), \mathbf{C}(\tau)\}), \forall \tau \in [t-L, t], \end{aligned} \quad (15)$$

where $\mathbf{p}(\tau)$ represents the ground-truth coordinates at time slot τ . By integrating these sources, the system utilizes RF observations to correct long-term sensor drift while leveraging sensing data to sustain trajectory estimation during signal interruptions, ensuring reliable localization in low-altitude environments.

C. Existing Approaches

Multimodal data fusion in ISAC systems has evolved through various paradigms. In this subsection, we categorize representative methods into two primary classes, namely static fusion strategies [25] and monolithic neural architectures [27]–[29], and highlight their respective limitations to motivate the design of our proposed method.

1) *Static fusion strategies*: The most straightforward approach to multimodal fusion relies on fixed rules or hand-crafted heuristics. Formally, let D denote the total number of modalities. For each modality d , an intermediate feature embedding $\mathbf{z}_d(t)$ is first extracted from the raw input. These modality-specific embeddings are then aggregated into a joint representation $\mathbf{z}(t)$ via deterministic operations:

- **Direct Concatenation** [25]: Feature vectors from different modalities are concatenated into a single high-dimensional vector

$$\mathbf{z}(t) = [\mathbf{z}_1(t)^T, \dots, \mathbf{z}_D(t)^T]^T, \quad (16)$$

which is subsequently fed into a downstream predictor.

- **Weighted Summation**: Features are combined using pre-assigned scalar weights as follows,

$$\mathbf{z}(t) = \sum_{d=1}^D w_d \mathbf{z}_d(t), \quad (17)$$

where w_d is typically determined by expert knowledge or heuristic confidence estimates.

Although computationally efficient and easy to implement, such static strategies fuse all inputs using pre-assigned and fixed weights, regardless of their instantaneous quality. This limitation is particularly critical in dynamic LAWNs, where sensor reliability fluctuates rapidly due to blockage or adverse environmental conditions. Consequently, static fusion strategies fail to filter out degraded sensing data, inevitably leading to performance degradation.

2) *Monolithic neural architectures*: To overcome the limitations of static fusion, recent research has shifted toward unified deep learning frameworks, such as multimodal transformers [27], [29] and MLLMs [28]. These approaches employ a monolithic architecture where heterogeneous inputs are projected onto a shared latent space and processed through a dense backbone to capture cross-modal correlations. However, this design presents critical bottlenecks for the practical ISAC deployment in LAWNs. First, sharing parameters across




Fig. 2. Architecture of the proposed multimodal MoE framework. Modality-specific experts extract features from heterogeneous inputs, which are aggregated via weighted summation using coefficients from a gating network to form a fused representation for the prediction head.

modalities often leads to modality interference, where dominant sensors (e.g., vision) suppress the learning of others, reducing robustness when the primary source fails. Second, these models enforce unconditional computation, activating the full network for every inference step regardless of data quality. This processing overhead creates a significant energy burden for battery-powered UAVs, hindering deployment in resource-constrained aerial scenarios.

In summary, static fusion strategies are computationally efficient but rely on fixed fusion rules that cannot adapt to changing modality quality or dynamic environments. In contrast, monolithic neural architectures that process all modalities jointly offer strong representation power but incur high computational cost and cross-modal interference because all modality-specific branches are always activated, even when some modalities are noisy or uninformative. To address both the lack of adaptivity in static fusion and the high cost of monolithic neural architectures, we propose a multimodal MoE framework in the following section. Instead of a single monolithic network, our approach uses a modular architecture in which specialized expert networks are coordinated by a gating mechanism that adaptively weights experts based on the current input. This design effectively combines environmental adaptability with computational efficiency, enabling the system to intelligently prioritize reliable sensing modalities for robust perception while minimizing the energy burden on resource-constrained aerial platforms.

III. PROPOSED MULTIMODAL MOE FRAMEWORK

In this section, we present the proposed multimodal MoE framework, specifically designed to address the dynamic requirements of ISAC in LAWNs. As illustrated in Fig. 2, the architecture consists of modality-specific expert networks and a modality-aware gating network. We first detail the

mathematical formulation of these components, followed by the description of the end-to-end training procedure.

A. Multimodal MoE Architecture

To effectively extract discriminative features from heterogeneous modalities, the proposed framework constructs a set of parallel expert networks. Unlike monolithic architectures that enforce a unified processing path, our design adopts a modular approach where each sensing modality is handled by a dedicated expert network. We define the d -th expert as a non-linear mapping function $g_d(\cdot; \theta_d)$, parameterized by the trainable weights θ_d . The complete ensemble of experts is denoted as $\mathcal{E} = \{g_d(\cdot; \theta_d)\}_{d=1}^D$. This decoupled structure allows each expert to employ a distinct inductive bias tailored to the specific data format of its corresponding modality. For a given time slot t , let $\mathbf{X}(t) = \{\mathbf{X}_d(t)\}_{d=1}^D$ represent the set of synchronized multimodal inputs. Each expert processes its specific input $\mathbf{X}_d(t)$ to generate the high-dimensional feature embedding $\mathbf{z}_d(t)$ as

$$\mathbf{z}_d(t) = g_d(\mathbf{X}_d(t); \theta_d), \forall d. \quad (18)$$

To improve the feature representation for downstream ISAC tasks, the neural architecture of each expert is specifically customized to align with the characteristics of its corresponding modality. For instance, since visual images are rich in spatial patterns and object details, the visual expert typically utilizes powerful feature extractors such as ResNet [31] or vision transformers [32] to interpret complex semantic context. In contrast, radar signals often contain time-varying motion information. To capture these dynamic attributes including target range and velocity, the radar expert employs architectures suited for sequential or spectral processing, such as convolutional neural networks [33] or long short-term memory models [34]. Furthermore, handling lidar data requires a different

approach because they consist of scattered and unordered three-dimensional points rather than a regular image grid. Consequently, the lidar expert leverages specialized geometric networks including PointNet [35] or voxel-based methods [36] to accurately reconstruct spatial depth and structure. This tailored design ensures that the unique strengths of each sensing modality are fully preserved and utilized in the final fusion process.

While independent feature extraction effectively captures modality-specific patterns, relying solely on isolated processing neglects the complex interactions between heterogeneous data modalities. In practical deployment over LAWNs, the reliability of each sensing modality fluctuates dynamically due to changing environmental conditions and sensor limitations. Consequently, fusion strategies that assign static or uniform weights often lead to suboptimal performance when some modalities are corrupted. To overcome this limitation, it is critical to equip the framework with an adaptive mechanism capable of assessing the instantaneous contribution of each expert. Motivated by this, we propose a modality-aware gating network that dynamically evaluates the relative importance of diverse inputs to generate context-aware fusion weights, as detailed in the subsequent subsection.

B. Modality-Aware Gating Network

The modality-aware gating network functions as the central decision module within the proposed multimodal MoE framework. While the expert networks focus on extracting deep semantic features from each modality, the gating network is tasked with evaluating the global context to determine the instantaneous reliability of each expert. In dynamic low-altitude environments, factors such as variable lighting, blockage, or signal clutter can cause the informativeness of individual modalities to fluctuate rapidly. Consequently, static fusion strategies that assign fixed weights fail to adapt to these variations and exhibit performance degradation when a primary sensor is compromised. To overcome this limitation, the gating network learns to model the uncertainty associated with each modality and dynamically generates a probability distribution over the experts, effectively prioritizing reliable signals while suppressing noisy or irrelevant inputs.

Formally, the gating network is defined as a learnable non-linear mapping $g_{\text{gate}}(\cdot; \theta_{\text{gate}})$ parameterized by θ_{gate} . It accepts the synchronized multimodal input set $\mathbf{X}(t)$ as context and computes a vector of normalized gating coefficients. These coefficients represent the relative contribution of each modality at time slot t and are calculated via the softmax function as

$$[w_1(t), w_2(t), \dots, w_D(t)] = \text{softmax}(g_{\text{gate}}(\mathbf{X}(t); \theta_{\text{gate}})), \quad (19)$$

where the softmax activation ensures that the generated weights are non-negative and sum to one, i.e.,

$$\begin{aligned} \sum_{d=1}^D w_d(t) &= 1, \\ w_d(t) &\geq 0, \forall d. \end{aligned} \quad (20)$$

In terms of architecture, the gating network is designed to be computationally lightweight to minimize inference latency.

It typically comprises a multilayer perceptron (MLP) with non-linear activation functions. The processing flow begins by aggregating the heterogeneous inputs into a shared latent embedding to capture cross-modal dependencies and environmental context. This compact representation is subsequently projected into a vector of raw logits, which are then normalized to produce the final fusion weights. Furthermore, the specific architecture of the gating network is adaptable to the complexity of the input data modality. For instance, when processing high-dimensional inputs like visual images or radar maps, the network may incorporate lightweight convolutional layers or global average pooling operations to extract global context descriptors before the dense gating layers. This flexible design enables accurate assessment of modality relevance without imposing a significant computational burden.

Upon obtaining the normalized gating weights, the final fused multimodal representation $\mathbf{z}(t)$ is computed as a weighted aggregation of the modality-specific expert embeddings $\{\mathbf{z}_d(t)\}_{d=1}^D$ as follows,

$$\mathbf{z}(t) = \sum_{d=1}^D w_d(t) \mathbf{z}_d(t). \quad (21)$$

This aggregated feature vector captures complementary information extracted from heterogeneous sensing modalities and serves as the input to the task-specific prediction head. Formally, we define the prediction head as a non-linear mapping $g_o(\cdot; \theta_o)$, parameterized by θ_o , which transforms the fused representation $\mathbf{z}(t)$ into the final task-specific prediction $\hat{\mathbf{o}}(t)$ as follows,

$$\hat{\mathbf{o}}(t) = g_o(\mathbf{z}(t); \theta_o). \quad (22)$$

The architecture of the prediction head $g_o(\cdot; \theta_o)$ is customized to the specific nature of the downstream ISAC task:

- For classification tasks, such as predicting the optimal beam index, the network ends with a softmax layer that outputs a probability distribution across all possible categories and is trained using a cross-entropy loss function to select the most likely candidates.
- For regression tasks, such as estimating channel parameters or tracking UAV trajectories, the network ends with linear output units. This configuration directly predicts continuous numerical values and is optimized by minimizing the error between the predicted and actual values, typically using MSE.

In summary, by jointly optimizing the modality-specific experts $\{g_d(\cdot; \theta_d)\}_{d=1}^D$, the adaptive gating network $g_{\text{gate}}(\cdot; \theta_{\text{gate}})$, and the prediction head $g_o(\cdot; \theta_o)$, the proposed multimodal MoE framework establishes a robust end-to-end learning paradigm for multimodal ISAC. This mechanism enables the system to dynamically balance modality contributions based on real-time environmental contexts, thereby significantly enhancing predictive accuracy and resilience against sensor uncertainties in dynamic low-altitude scenarios.

C. Algorithm Development

In this subsection, we present the training procedure for the proposed multimodal MoE framework in detail. Specifically,

our objective is to jointly optimize the parameters of the expert networks, the modality-aware gating network, and the subsequent prediction head. To this end, we formulate the training as a general supervised learning problem using a labeled multimodal dataset.

Given a training dataset of synchronized multimodal inputs and ground-truth labels, denoted as $\mathcal{D} = \{(\mathbf{X}(t), \mathbf{o}(t))\}_{t=1}^T$, our goal is to minimize the empirical risk as measured by the loss function $\mathcal{L}(\cdot)$. Depending on the specific ISAC task, the loss function $\mathcal{L}(\cdot)$ can be chosen as binary cross-entropy (BCE) for classification tasks (e.g., proactive blockage prediction), MSE for regression tasks (e.g., channel estimation), or negative sum-rate loss for beamforming optimization tasks. Formally, the complete parameter optimization problem can be expressed as

$$\min_{\{\theta_d\}_{d=1}^D, \theta_{\text{gate}}, \theta_o} \frac{1}{T} \sum_{t=1}^T \mathcal{L}(\hat{\mathbf{o}}(t), \mathbf{o}(t)), \quad (23)$$

where the prediction output $\hat{\mathbf{o}}(t)$ is obtained by sequentially computing expert outputs, modality-aware fusion weights, fused representations, and final predictions as follows,

$$\begin{aligned} \mathbf{z}_d(t) &= g_d(\mathbf{X}_d(t); \theta_d), \quad \forall d, \\ [w_1(t), \dots, w_D(t)] &= \text{softmax}(g_{\text{gate}}(\mathbf{X}(t); \theta_{\text{gate}})), \\ \mathbf{z}(t) &= \sum_{d=1}^D w_d(t) \mathbf{z}_d(t), \\ \hat{\mathbf{o}}(t) &= g_o(\mathbf{z}(t); \theta_o). \end{aligned} \quad (24)$$

The complete end-to-end training procedure is formally summarized in Algorithm 1. By minimizing the task-specific loss $\mathcal{L}(\cdot)$ via backpropagation, the framework jointly optimizes the expert backbones, the gating mechanism, and the prediction head within a unified computational graph. This joint training process allows the model to learn a context-aware fusion policy that dynamically recalibrates modality weights based on the instantaneous environmental state. As a result, the trained multimodal MoE framework tends to exploit information from reliable sensing modalities while reducing the influence of degraded ones, leading to robust performance across diverse ISAC applications in dynamic low-altitude environments.

IV. SPARSE MULTIMODAL MOE FRAMEWORK

In this section, we propose a sparse multimodal MoE framework to meet the stringent energy and computation constraints of the aerial platforms in LAWNs. By selectively activating only the top- N most relevant experts via a sparse routing mechanism, this variant avoids unnecessary processing and reduces resource overhead while maintaining adaptive fusion.

A. Sparse Expert Activation

While the proposed multimodal MoE framework offers comprehensive fusion, activating all D experts for every inference step imposes a computational burden that challenges the energy constraints of aerial platforms. However, as evidenced by the performance comparison in Fig. 3, the accuracy gain

Algorithm 1: Training Procedure of the Proposed Multimodal MoE Framework.

Input: Training dataset $\mathcal{D} = \{(\mathbf{X}(t), \mathbf{o}(t))\}_{t=1}^T$, initialized expert network parameters $\{\theta_d\}_{d=1}^D$, gating network parameters θ_{gate} , prediction head parameters θ_o , learning rate η .

- 1 **for** $epoch = 1, 2, \dots, E$ **do**
- 2 Shuffle the training dataset $\mathcal{D}$;
- 3 **for** *each training sample* $(\mathbf{X}(t), \mathbf{o}(t)) \in \mathcal{D}$ **do**
- 4 Compute modality-specific expert embeddings:

$$\mathbf{z}_d(t) = g_d(\mathbf{X}_d(t); \theta_d), \quad \forall d; \quad (25)$$
- 5 Compute modality-aware fusion weights:

$$[w_1(t), \dots, w_D(t)] = \text{softmax}(g_{\text{gate}}(\mathbf{X}(t); \theta_{\text{gate}})); \quad (26)$$
- 6 Compute fused multimodal representation:

$$\mathbf{z}(t) = \sum_{d=1}^D w_d(t) \mathbf{z}_d(t); \quad (27)$$
- 7 Compute final prediction output:

$$\hat{\mathbf{o}}(t) = g_o(\mathbf{z}(t); \theta_o); \quad (28)$$
- 8 Evaluate supervised loss function:

$$\mathcal{L}_t = \mathcal{L}(\mathbf{o}(t), \hat{\mathbf{o}}(t)); \quad (29)$$
- 9 Compute gradients via backpropagation:

$$\nabla_{\theta_d} \mathcal{L}_t, \forall d; \quad \nabla_{\theta_{\text{gate}}} \mathcal{L}_t; \quad \nabla_{\theta_o} \mathcal{L}_t; \quad (30)$$
- 10 Update parameters with gradient descent:

$$\begin{aligned} \theta_d &\leftarrow \theta_d - \eta \nabla_{\theta_d} \mathcal{L}_t, \quad \forall d; \\ \theta_{\text{gate}} &\leftarrow \theta_{\text{gate}} - \eta \nabla_{\theta_{\text{gate}}} \mathcal{L}_t; \\ \theta_o &\leftarrow \theta_o - \eta \nabla_{\theta_o} \mathcal{L}_t; \end{aligned} \quad (31)$$
- 11 **end**
- 12 **end**

Output: Expert parameters $\{\theta_d\}_{d=1}^D$, gating parameter θ_{gate} , and prediction head parameter θ_o .

from integrating additional modalities exhibits clear diminishing returns. Specifically, the marginal improvement decreases significantly as the number of modalities grows (e.g., the gap between 3 and 4 modalities is minimal), implying that fully activating the entire expert ensemble is often unnecessary and computationally redundant.

To capitalize on this observation, we propose a sparse multimodal MoE framework that adaptively activates different experts. By dynamically selecting only the top- N most informative experts per input (with $N < D$), this approach effectively prunes redundant branches. This yields a favorable trade-off between inference accuracy and computational efficiency for real-time aerial applications. Specifically, the sparse MoE shares the same expert pool and gating network as the dense model. At each time slot t , the gating network first produces a set of fusion scores $[w_1(t), \dots, w_D(t)]$ as in the dense case. Instead of using all experts, the sparse variant only activates those associated with the largest scores. Concretely,




Fig. 3. Top-1 beam prediction accuracy across different combinations of sensing modalities.

we select the indices of the top- N experts that are most relevant for the current multimodal observation as

$$\mathcal{S}(t) = \text{TopN}(w_1(t), \dots, w_D(t)), \quad (32)$$

where $\text{TopN}(\cdot)$ returns the indices of the N largest entries among $\{w_d(t)\}_{d=1}^D$. Based on this set, we define a binary activation mask to indicate whether the d -th expert is activated at time slot t , given by

$$m_d(t) = \begin{cases} 1, & \text{if } d \in \mathcal{S}(t), \\ 0, & \text{otherwise,} \end{cases} \quad d = 1, \dots, D.$$

In the multimodal MoE proposed in Section III, all expert embeddings $\{\mathbf{z}_d(t)\}_{d=1}^D$ are computed and combined according to the normalized weights. In contrast, the sparse MoE only evaluates the experts with $m_d(t) = 1$. That is, for each time slot t , we compute

$$\mathbf{z}_d(t) = g_d(\mathbf{X}_d(t); \boldsymbol{\theta}_d) \quad \text{only if } m_d(t) = 1,$$

and skip the forward pass through the remaining experts, which directly leads to a reduction in computational cost and memory access. To ensure that the fused representation remains a convex combination of the active expert outputs, we renormalize the original weights over the selected subset as

$$\tilde{w}_d(t) = \frac{m_d(t)w_d(t)}{\sum_{j=1}^D m_j(t)w_j(t) + \varepsilon}, \quad d = 1, \dots, D, \quad (33)$$

where $\varepsilon > 0$ is a small constant added for numerical stability. By construction, $\tilde{w}_d(t) \geq 0$ and $\sum_{d=1}^D \tilde{w}_d(t) = 1$, and $\tilde{w}_d(t) = 0$ for all inactive experts with $m_d(t) = 0$. With these sparsified weights, the fused multimodal representation is given by

$$\tilde{\mathbf{z}}(t) = \sum_{d=1}^D \tilde{w}_d(t) \mathbf{z}_d(t), \quad (34)$$

where the sum effectively runs only over the N active experts. Compared to the dense case, the sparse MoE preserves the gating behavior, since larger $w_d(t)$ still leads to larger normalized weights $\tilde{w}_d(t)$, while reducing computation and memory cost from evaluating all D experts to only $N < D$ experts per time slot.

B. Training Procedure with Sparse Activation

The integration of the top- N selection operator introduces a non-differentiable step into the computational graph, which blocks the flow of gradients required for standard back-propagation. To circumvent this issue and enable end-to-end optimization, we employ the straight-through estimator (STE) [37]. This technique decouples the forward execution from the backward gradient approximation, allowing the network to make hard routing decisions while remaining differentiable for parameter updates.

During the forward propagation phase, the system performs selective execution, activating only the experts deemed most relevant. The gating network generates the raw probability vector $\mathbf{w}(t)$, which is then processed by the non-differentiable top- N operator to yield a binary mask and the renormalized sparse weights $\tilde{\mathbf{w}}(t)$. Consequently, only the selected subset of experts contributes to the fused representation $\mathbf{z}(t)$, improving computational efficiency. In contrast, during the backward propagation phase, the STE heuristic approximates the hard top- N operation as an identity function for the purpose of gradient calculation. Specifically, we assume that the gradients flow through the discrete selection layer without attenuation. Formally, the gradient of the loss function $\mathcal{L}$ with respect to the original gating weights $\mathbf{w}(t)$ is approximated by the gradients computed for the sparsified weights $\tilde{\mathbf{w}}(t)$ as follows,

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}(t)} \approx \frac{\partial \mathcal{L}}{\partial \tilde{\mathbf{w}}(t)}. \quad (35)$$

This approximation implies that if an expert is activated, its gradient contribution is backpropagated directly to the gating network. If an expert is inactive (i.e., $\tilde{w}_d(t) = 0$), it receives no gradient update for that instance. This mechanism allows the gating parameter $\boldsymbol{\theta}_g$ to be updated via stochastic gradient descent despite the discrete nature of the inference process. Under this training paradigm, the overall learning objective remains consistent with the empirical risk minimization problem in (23). By leveraging the STE, the sparse MoE framework effectively learns an adaptive and sparsity-constrained routing policy that balances prediction accuracy with the stringent resource limitations in practical LAWNs.

C. Computational Benefits and Practical Considerations

The sparse MoE framework inherits the architectural flexibility of the multimodal MoE in Section III, while achieving significantly improved computational efficiency. Given that only N experts are evaluated per input (with $N < D$), the total computational cost per forward pass is reduced from $\mathcal{O}(D)$ to $\mathcal{O}(N)$. The same reduction applies to memory and energy usage, making the sparse MoE model particularly attractive for deployment on embedded processors and real-time UAV platforms. Additionally, by enforcing sparsity at inference time, the model inherently introduces a form of regularization that mitigates overfitting, especially under limited training data or noisy sensor inputs. The ability to deactivate uninformative or redundant modalities also improves robustness under partial modality failures, particularly when vision or lidar becomes unreliable during adverse weather conditions.

In practice, the value of N can be treated as a tunable hyperparameter that balances accuracy and efficiency. A smaller N yields faster inference but may omit useful modality information, while a larger N improves performance but increases resource demands. To enhance stability during training, it is often beneficial to start with a relatively larger N and gradually anneal it using a predefined schedule or entropy-based gating regularization, so that the model progressively learns to rely on a compact subset of experts. Moreover, the sparse MoE framework is fully compatible with other multimodal extensions, including hierarchical expert routing [38] and uncertainty-aware gating [39]. These extensions can further improve specialization and generalization in complex ISAC scenarios, and can be readily integrated on top of the sparse routing mechanism.

In summary, the sparse multimodal MoE framework provides a lightweight yet adaptive solution for efficient ISAC inference. By activating only a subset of experts based on modality-aware gating, the model achieves favorable trade-offs between computational cost, robustness, and prediction performance. In the next section, we evaluate the proposed multimodal MoE and sparse multimodal MoE frameworks under multiple ISAC tasks in realistic LAWN scenarios to quantitatively demonstrate these benefits.


V. SIMULATION RESULTS

In this section, we conduct comprehensive simulations to evaluate the effectiveness of the proposed multimodal MoE frameworks in realistic LAWN scenarios. The evaluation covers three representative ISAC tasks under various baseline fusion strategies and performance metrics.


A. Simulation Setups

We evaluate the proposed multimodal MoE frameworks using a public multimodal ISAC dataset for low-altitude scenarios [40], [41], which provides synchronized RF communication data and heterogeneous sensing modalities. Fig. 4 illustrates a representative scene from the dataset, showing the environment from both the ground BS and UAV perspectives, with UAVs highlighted in red boxes. We consider three representative ISAC tasks: i) *sensing-aided beam prediction*, where the model predicts the optimal beam index from sensing inputs and is evaluated by top-1 beam accuracy; ii) *sensing-aided path loss prediction*, where the model predicts downlink path loss and is evaluated by the MSE between estimated and ground-truth values; and iii) *communication-aided UAV trajectory tracking*, where the UAV position is regressed from sensing and communication features, and evaluated by MSE and average Euclidean distance between predicted and ground-truth positions.

For the sensing model, each UAV is equipped with an RGB camera and lidar sensor, while the BS carries frequency-modulated continuous wave (FMCW) radars operating at 77–81 GHz. The sensing data includes RGB images, lidar point clouds, mmWave radar point clouds, and GPS position information captured across 3,000 time slots. For the communication model, we simulate the wireless channels at



(a) UAV perspective.



(b) Ground BS perspective.

Fig. 4. Illustration of a representative scene from the multimodal ISAC dataset in LAWNs, showing the environment from the ground BS and UAV perspectives. UAVs are marked in red boxes.

5.915 GHz with 20 MHz bandwidth. The UAVs are equipped with 4-element ULAs, and the resulting data includes 30,000 time slots of complex-valued channel impulse responses and path loss values under dynamic propagation conditions. The RF communication and multimodal sensing data are fully synchronized to enable multimodal fusion.

For the learning model, we implement a multimodal MoE architecture with 5 sensing modalities: vision images, lidar, radar, position information, and RF communication signals. Each modality is assigned a group of 3 parallel experts, resulting in a total of 15 expert subnetworks. In the sparse MoE variant, we activate a subset of 5 experts out of the total 15 for each input, and their weights are renormalized prior to aggregation. For image-based modalities, each expert adopts a ResNet-18 backbone [31]; for point-cloud-based modalities, each expert is implemented using a PointNet backbone [35]. The modality-aware gating network is a three-layer MLP with 128 hidden units and ReLU activations, followed by a softmax layer to produce normalized fusion weights. All models are trained using the Adam optimizer with a learning rate of 10^{-3} and batch size of 64. We split the entire dataset into 70% for training and 30% for testing, and all results are reported on the testing set.

To evaluate the effectiveness of the proposed multimodal




Fig. 5. Comparison of Top-1 beam prediction accuracy for the proposed multimodal MoE frameworks and benchmark methods.

MoE and sparse MoE frameworks, we compare them against the following representative baselines.

- **Single-modality baselines** [42], [43]: We include vision-only and lidar-only models, where each modality is processed independently using the same expert network architecture without any fusion. These serve as lower-bound references for unimodal performance.
- **Feature concatenation fusion** [25]: Multimodal features from all modalities are concatenated after backbone encoding and fed into a shared MLP prediction head. This represents a common and straightforward fusion strategy.
- **Multimodal transformer** [27], [29]: A transformer encoder is employed to model cross-modal interactions by attending over modality-specific token embeddings. Self-attention and cross-attention mechanisms enable the network to dynamically capture inter-modal dependencies and contextual correlations across sensing streams.
- **Multimodal LLM** [28]: We adopt a pretrained LLM by projecting sensor features into a shared embedding space followed by fine-tuning. This baseline highlights the performance of large-scale pretrained fusion.

B. Sensing-Aided Beam Prediction

In this subsection, we evaluate model performance on the sensing-aided beam prediction task, where the goal is to select the optimal downlink beam index based on multimodal sensing inputs. As shown in Fig. 5, the proposed multimodal MoE and sparse MoE frameworks significantly outperform all baselines in terms of top-1 prediction accuracy. The dense multimodal MoE model achieves the best final performance, while the sparse MoE maintains competitive accuracy with fewer active experts, demonstrating its efficiency. The superior results can be attributed to the adaptive gating mechanism, which assigns dynamic fusion weights to expert subnetworks, allowing the model to selectively emphasize informative modalities under varying channel and environmental conditions.




Fig. 6. Comparison of NMSE performance for sensing-aided path loss prediction using the proposed multimodal MoE frameworks and benchmark methods.

Among the baselines, the multimodal transformer and multimodal LLM benefit from stronger representation capacity and attention mechanisms, achieving better performance than early fusion via feature concatenation. However, their performance remains below the MoE-based models due to the lack of explicit expert specialization across modalities. Unimodal baselines (vision-only and lidar-only) show limited accuracy, confirming that leveraging diverse and complementary sensing modalities is essential for robust beam prediction in dynamic low-altitude environments.

C. Sensing-Aided Path Loss Prediction

In this subsection, we evaluate all methods on the task of sensing-aided path loss prediction, where the goal is to estimate the downlink path loss using multimodal sensing features. Fig. 6 reports the normalized mean squared error (NMSE) curves in dB scale over training epochs. The proposed multimodal MoE achieves the lowest NMSE, followed closely by the sparse MoE, which performs nearly on par while requiring fewer active experts. These results highlight the advantage of using adaptive modality-aware expert selection to adaptively predict path loss under varying environmental conditions and blockage states.

Compared to the baselines, the multimodal transformer and LLM deliver improved performance over early fusion approaches by modeling cross-modal correlations and utilizing large-scale pretraining, respectively. However, they remain outperformed by MoE-based methods due to their lack of explicit expert specialization and adaptive routing. Meanwhile, unimodal baselines (vision-only and lidar-only) suffer from limited expressiveness, reflecting the insufficiency of single-modality inputs in accurately capturing fine-grained propagation loss variations in complex environments.





Fig. 7. Comparison of NMSE performance for communication-aided UAV trajectory tracking using the multimodal MoE frameworks and benchmark methods.

D. Communication-Aided UAV Trajectory Tracking


In this subsection, we evaluate the proposed multimodal MoE method and sparse MoE on the task of UAV trajectory tracking, where the objective is to predict the UAV's 3D position at each time step based on multimodal sensing and RF communication inputs. This task is essential for autonomous navigation in dynamic low-altitude environments, where accurate trajectory estimation enables proactive control and reliable mobility support.

Fig. 7 shows the NMSE in dB scale across training epochs. As observed, the proposed multimodal MoE achieves the lowest NMSE, indicating superior estimation accuracy, while the sparse MoE maintains comparable performance with reduced computational overhead. These results validate the effectiveness of adaptive expert selection in capturing complex path-dependent signal variations. Among the baselines, the multimodal transformer and multimodal LLM achieve moderate improvements over the static fusion baseline by leveraging attention mechanisms and pretrained multimodal embeddings, respectively. However, their fusion lacks explicit modality specialization, resulting in inferior performance compared to the MoE-based architectures. Unimodal methods (vision-only and lidar-only) consistently underperform, highlighting the limitations of single-modality inputs in accurately reconstructing continuous trajectories amidst sensor drift and environmental blockage.

To further illustrate the model's performance, Fig. 8a and Fig. 8b visualize the predicted trajectory generated by the multimodal MoE alongside the ground-truth UAV path. The predicted positions exhibit close alignment with the true trajectory, particularly around direction changes and elevation variations, demonstrating the model's ability to accurately track UAV motion in dynamic low-altitude scenarios.



(a) Scenario 1



(b) Scenario 2

Fig. 8. Comparison between predicted and ground-truth UAV trajectories using the proposed multimodal MoE framework.

VI. CONCLUSION

In this paper, we have investigated multimodal ISAC in LAWNs, leveraging heterogeneous sensing modalities and communication signals to achieve robust environmental perception and reliable connectivity in dynamic aerial environments. We have formulated a general learning-based optimization framework for multimodal ISAC, which has been instantiated across three representative ISAC tasks. To address the challenges of modality heterogeneity and time-varying reliability, we have proposed a multimodal MoE architecture that adaptively assigns modality-specific fusion weights via a learnable gating network. Additionally, we have developed a sparse MoE variant that reduces inference costs by selectively activating a subset of expert networks, while preserving the benefits of adaptive fusion. Extensive simulations on three representative ISAC tasks in LAWNs have demonstrated that the proposed methods consistently outperform conventional multimodal fusion baselines in terms of learning performance and training sample efficiency.

REFERENCES

- [1] K. Zhang, W. Yu, H. He, S. Song, J. Zhang, and K. B. Letaief, "Multimodal deep learning-empowered beam prediction in future THz ISAC systems," in *Proc. IEEE Int. Symp. Pers., Indoor Mobile Radio Commun. (PIMRC)*, Istanbul, Turkey, pp. 1–6, Sept. 2025.
- [2] K. B. Letaief, W. Chen, Y. Shi, J. Zhang, and Y.-J. A. Zhang, "The roadmap to 6G: AI empowered wireless networks," *IEEE Commun. Mag.*, vol. 57, no. 8, pp. 84–90, 2019.
- [3] K. B. Letaief, Y. Shi, J. Lu, and J. Lu, "Edge artificial intelligence for 6G: Vision, enabling technologies, and applications," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 1, pp. 5–36, 2022.
- [4] J. Qiu, D. Grace, G. Ding, M. D. Zakaria, and Q. Wu, "Air-ground heterogeneous networks for 5G and beyond via integrating high and low altitude platforms," *IEEE Wireless Commun.*, vol. 26, no. 6, pp. 140–148, 2019.
- [5] W. Yuan, Y. Cui, J. Wang, F. Liu, G. Sun, T. Xiang, J. Xu, S. Jin, D. Niyato, S. Coleri, S. Sun, S. Mao, A. Jamalipour, D. I. Kim, M.-S. Alouini, and X. Shen, "From ground to sky: Architectures, applications, and challenges shaping low-altitude wireless networks," *arXiv preprint arXiv:2506.12308*, 2025.
- [6] Y. Lin, T. Wang, and S. Wang, "UAV-assisted emergency communications: An extended multi-armed bandit perspective," *IEEE Commun. Lett.*, vol. 23, no. 5, pp. 938–941, 2019.
- [7] K. Kuru, D. Ansell, W. Khan, and H. Yetgin, "Analysis and optimization of unmanned aerial vehicle swarms in logistics: An intelligent delivery platform," *IEEE Access*, vol. 7, pp. 15804–15831, 2019.
- [8] C. Zhang and J. M. Kovacs, "The application of small unmanned aerial systems for precision agriculture: a review," *Precision Agriculture*, vol. 13, no. 6, pp. 693–712, 2012.
- [9] M. Zhang and X. Li, "Drone-enabled internet-of-things relay for environmental monitoring in remote areas without public networks," *IEEE Internet Things J.*, vol. 7, no. 8, pp. 7648–7662, 2020.
- [10] Y. Zhu and S. Wang, "Efficient aerial data collection with cooperative trajectory planning for large-scale wireless sensor networks," *IEEE Trans. Commun.*, vol. 70, no. 1, pp. 433–444, 2022.
- [11] F. Liu, Y. Cui, C. Masouros, J. Xu, T. X. Han, Y. C. Eldar, and S. Buzzi, "Integrated sensing and communications: Toward dual-functional wireless networks for 6G and beyond," *IEEE J. Sel. Areas Commun.*, vol. 40, no. 6, pp. 1728–1767, 2022.
- [12] J. A. Zhang, F. Liu, C. Masouros, R. W. Heath, Z. Feng, L. Zheng, and A. Petropulu, "An overview of signal processing techniques for joint communication and radar sensing," *IEEE J. Sel. Top. Signal Process.*, vol. 15, no. 6, pp. 1295–1315, 2021.
- [13] Z. Lyu, G. Zhu, and J. Xu, "Joint maneuver and beamforming design for UAV-enabled integrated sensing and communication," *IEEE Trans. Wireless Commun.*, vol. 22, no. 4, pp. 2424–2440, 2023.
- [14] K. Meng, Q. Wu, S. Ma, W. Chen, K. Wang, and J. Li, "Throughput maximization for UAV-enabled integrated periodic sensing and communication," *IEEE Trans. Wireless Commun.*, vol. 22, no. 1, pp. 671–687, 2023.
- [15] Z. Wu, X. Li, Y. Cai, and W. Yuan, "Joint trajectory and resource allocation design for RIS-assisted UAV-enabled ISAC systems," *IEEE Wireless Commun. Lett.*, vol. 13, no. 5, pp. 1384–1388, 2024.
- [16] S. Zhou, H. Yang, L. Xiang, and K. Yang, "Temporal-assisted beamforming and trajectory prediction in sensing-enabled UAV communications," *IEEE Trans. Commun.*, vol. 73, no. 7, pp. 5408–5419, 2025.
- [17] Y. Jiang, Q. Wu, W. Chen, and K. Meng, "UAV-enabled integrated sensing and communication: Tracking design and optimization," *IEEE Commun. Lett.*, vol. 28, no. 5, pp. 1024–1028, 2024.
- [18] Y. Jiang, Q. Wu, H. Hui, W. Chen, and D. W. K. Ng, "Low-altitude uav tracking via sensing-assisted predictive beamforming," *arXiv preprint arXiv:2509.12698*, 2025.
- [19] K. Meng, Q. Wu, W. Chen, and D. Li, "Sensing-assisted communication in vehicular networks with intelligent surface," *IEEE Trans. Veh. Technol.*, vol. 73, no. 1, pp. 876–893, 2024.
- [20] C. Ji, Q. Liu, J. Liu, C. Yu, Y. Sun, R. Wang, and F. Liu, "Doppler-based multistatic drone tracking via cellular downlink signals," *arXiv preprint arXiv:2509.25732*, 2025.
- [21] Y. Yang, Z. Li, J. Zhao, Q. Wu, Z. Wei, W. Chen, and W. Jia, "Cooperative multi-static ISAC networks: A unified design framework for active and passive sensing," *IEEE Trans. Wireless Commun.*, pp. 1–16, 2025.
- [22] X. Lu, Z. Wei, R. Xu, L. Wang, B. Lu, and J. Piao, "Integrated sensing and communication enabled multiple base stations cooperative UAV detection," in *Proc. IEEE Int. Conf. Commun. Workshops (ICC Workshops)*, pp. 1882–1887, 2024.
- [23] X. Cheng, H. Zhang, J. Zhang, S. Gao, S. Li, Z. Huang, L. Bai, Z. Yang, X. Zheng, and L. Yang, "Intelligent multi-modal sensing-communication integration: Synesthesia of machines," *IEEE Commun. Surveys Tuts.*, vol. 26, no. 1, pp. 258–301, 2024.
- [24] K. Tan and C. Zhu, "Multimodal sensing for intelligent V2X: A review of recent advances toward deployment," *IEEE Internet Things J.*, vol. 12, no. 22, pp. 46294–46315, 2025.
- [25] Q. Zhu, Y. Wang, W. Li, H. Huang, and G. Gui, "Advancing multi-modal beam prediction with cross-modal feature enhancement and dynamic fusion mechanism," *IEEE Trans. Commun.*, vol. 73, no. 9, pp. 7931–7940, 2025.
- [26] S. Tariq, B. E. Arfeto, U. Khalid, S. Kim, T. Q. Duong, and H. Shin, "Deep quantum-transformer networks for multimodal beam prediction in ISAC systems," *IEEE Internet Things J.*, vol. 11, no. 18, pp. 29387–29401, 2024.
- [27] Y. Cui, J. Nie, X. Cao, T. Yu, J. Zou, J. Mu, and X. Jing, "Sensing-assisted high reliable communication: A transformer-based beamforming approach," *IEEE J. Sel. Top. Signal Process.*, vol. 18, no. 5, pp. 782–795, 2024.
- [28] L. Cheng, H. Zhang, B. Di, D. Niyato, and L. Song, "Large language models empower multimodal integrated sensing and communication," *IEEE Commun. Mag.*, vol. 63, no. 5, pp. 190–197, 2025.
- [29] Y. Xie, Y. Gong, Y. Zhao, M. Zhou, S. Wang, D. Zhang, Y. Wang, and J. Wang, "Mpfusionnet: Transformer-based multi-modal perception fusion for predictive beamforming in low-altitude UAV communication networks," *IEEE Internet Things J.*, pp. 1–16, 2025.
- [30] M. Alrabeiah and A. Alkhateeb, "Deep learning for TDD and FDD massive MIMO: Mapping channels in space and frequency," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, pp. 1465–1470, 2019.
- [31] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 770–778, 2016.
- [32] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Representations (ICLR)*, 2021.
- [33] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol. 25, 2012.
- [34] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [35] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp. 652–660, 2017.
- [36] J. Deng, S. Shi, P. Li, W. Zhou, Y. Zhang, and H. Li, "Voxel r-cnn: Towards high performance voxel-based 3d object detection," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, pp. 1201–1209, 2021.
- [37] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *arXiv preprint arXiv:1308.3432*, 2013.
- [38] W. Li, D. Wang, Z. Ding, A. Sohrabizadeh, Z. Qin, J. Cong, and Y. Sun, "Hierarchical mixture of experts: Generalizable learning for high-level synthesis," in *Proc. AAAI Conf. Artif. Intell.*, vol. 39, pp. 18476–18484, 2025.
- [39] R. Zhang, Y. Luo, J. Liu, H. Yang, Z. Dong, D. Gudovskiy, T. Okuno, Y. Nakata, K. Keutzer, Y. Du, et al., "Efficient deweather mixture-of-experts with uncertainty-aware feature-wise linear modulation," in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, pp. 16812–16820, 2024.
- [40] X. Cheng, Z. Huang, Y. Yu, L. Bai, M. Sun, Z. Han, R. Zhang, and S. Li, "SynthSoM: A synthetic intelligent multi-modal sensing-communication dataset for synesthesia of machines (som)," *Scientific Data*, vol. 12, no. 1, p. 819, 2025.
- [41] X. Cheng, Z. Huang, L. Bai, H. Zhang, M. Sun, B. Liu, S. Li, J. Zhang, and M. Lee, "M3SC: A generic dataset for mixed multi-modal (MMM) sensing and communication integration," *China Communications*, vol. 20, no. 11, pp. 13–29, 2023.
- [42] M. Alrabeiah, A. Hredzak, and A. Alkhateeb, "Millimeter wave base stations with cameras: Vision-aided beam and blockage prediction," in *Proc. IEEE Veh. Technol. Conf. (VTC)*, pp. 1–5, 2020.
- [43] S. Jiang, G. Charan, and A. Alkhateeb, "Lidar aided future beam prediction in real-world millimeter wave V2I communications," *IEEE Wireless Commun. Lett.*, vol. 12, no. 2, pp. 212–216, 2023.