

Review

Large Model Enabled Embodied Intelligence for 6G Integrated Perception, Communication, and Computation Network

Zhuoran Li^{1,2}, Zhen Gao^{3,4,5,6}, Xinhua Liu^{1,2}, Zheng Wang⁷, Xiaotian Zhou^{8,9}, Lei Liu¹⁰, Yongpeng Wu¹¹, Wei Feng¹², Yongming Huang^{13,14}

1. School of Information and Electronics, Beijing Institute of Technology, Beijing 100081, China
2. School of Interdisciplinary Science, Beijing Institute of Technology, Beijing 100081, China
3. State Key Laboratory of CNS/ATM and the MIT Key Laboratory of Complex-Field Intelligent Sensing, Beijing 100081, China
4. Beijing Institute of Technology (BIT), BIT, Zhuhai 519088, China
5. Advanced Technology Research Institute, BIT (Jinan), Jinan 250307, China
6. Yangtze Delta Region Academy, BIT (Jiaxing), Jiaxing 314019, China
7. School of School of Information Science and Engineering, Southeast University, Nanjing 210096, China
8. Institute of Intelligent Communication Technologies, Shandong University, Jinan 250061, China
9. Shandong Key Laboratory of Intelligent Communication and Sensing-Computing Integration Shandong University, Jinan 250061, China
10. Zhejiang Provincial Key Laboratory of Information Processing, Communication and Networking, College of Information Science and Electronic Engineering, Zhejiang University, Hangzhou 310007, China
11. Department of Electronic Engineering, Shanghai Jiao Tong University, Minhang 200240, China
12. Department of Electronic Engineering, State Key Laboratory of Space Network and Communications, Tsinghua University, Beijing 100084, China
13. National Mobile Communications Research Laboratory, School of Information Science and Engineering, Southeast University, Nanjing 210096, China
14. Purple Mountain Laboratories, Nanjing 211111, China

Corresponding author: Zhen Gao; Email: gaozhen16@bit.edu.cn.

Received x xx, xxxx; Accepted x xx, xxx; Published x xx, xxxx.

Abstract—The advent of sixth-generation (6G) places intelligence at the core of wireless architecture, fusing perception, communication, and computation into a single closed-loop. This paper argues that large artificial intelligence models (LAMs) can endow base stations with perception, reasoning, and acting capabilities, thus transforming them into intelligent base station agents (IBSAs). We first review the historical evolution of BSs from single-functional analog infrastructure to distributed, software-defined, and finally LAM-empowered IBSA, highlighting the accompanying changes in architecture, hardware platforms, and deployment. We then present an IBSA architecture that couples a perception-cognition-execution pipeline with cloud-edge-end collaboration and parameter-efficient adaptation. Subsequently, we study two representative scenarios: (i) cooperative vehicle-road perception for autonomous driving, and (ii) ubiquitous base station support for low-altitude uncrewed aerial vehicle safety monitoring and response against unauthorized drones. On this basis, we analyze key enabling technologies spanning LAM design and training, efficient edge-cloud inference, multi-modal perception and actuation, as well as trustworthy security and governance. We further propose a holistic evaluation framework and benchmark considerations that jointly cover communication performance, perception accuracy, decision-making reliability, safety, and energy efficiency. Finally, we distill open challenges on benchmarks, continual adaptation, trustworthy decision-making, and standardization. Together, this work positions LAM-enabled IBSAs as a practical path toward integrated perception, communication, and computation native, safety-critical 6G systems.

Keywords—6G, integrated perception, communication, and computation, intelligent base station agent (IBSA), large AI model (LAM), multimodal fusion.

I. Introduction

1. Background and Motivation

In the sixth-generation (6G) era, intelligence is no longer an afterthought but a native feature woven into every network layer. Unlike past generations, 6G elevates artificial intelligence

(AI) from a tool to a fundamental infrastructure component, enabling networks to learn, adapt, and self-optimize in real time [1]. This AI-native paradigm, often termed “native AI” or “AI-native wireless systems”, envisions embedding machine-learning models directly into the radio inter-

face and control planes to achieve ultra-low latency, high reliability, and context-aware decision-making [2]. Recent work on artificial general intelligence (AGI)-native networks argues that simply scaling large language or generative models is insufficient; true network autonomy requires modules for perception, causal world modeling, and action planning to handle unforeseen scenarios with common-sense reasoning [2]. In this vision, large language models (LLMs) translate operator intents into network objectives, while specialized foundation models predict environmental dynamics and proactively adjust resource allocation. A quintessential application of native AI in 6G is integrated sensing and communications (ISAC), where sensing and payload transmission share spectrum, hardware, and intelligence pipelines [3]. By co-designing radar-like perception with adaptive beamforming, ISAC exemplifies how native AI can unlock mutual performance gains, turning base stations (BSs) into perceptive agents that both see and serve the environment [4].

2. Necessity of Large Models

The evolution of 6G networks necessitates a transition from static resource management to dynamic, intelligent adaptation. Foundational frameworks for ISAC, such as those comprehensively reviewed in [4], have successfully established the theoretical limits and signal processing principles for dual-functional networks. However, these conventional approaches are predominantly model-driven and rule-based. They rely on explicit mathematical modeling of the channel and environment, which often becomes intractable or inaccurate in highly dynamic, non-stationary, and semantically complex real-world scenarios. While traditional AI has been deployed to address these model deficits and handle non-linear complexities, these traditional AI systems are often designed as highly modular and single-function “silos,” focusing on optimization within specific domains, such as pure image recognition, natural language processing, or network routing. This fragmented intelligent architecture exhibits inherent limitations when facing complex, dynamic, and cross-modal challenges in the real world. With the advancement of the AGI paradigm, the strategic objective of the scientific community has shifted toward building unified systems that feature integrated perception, communication, and computation[5–7].

The integrated perception, communication, and computation paradigm aims to dismantle traditional functional boundaries, leveraging ubiquitous BS infrastructure to drive its transformation from a single-function communication node into an intelligent hub. This evolution represents not merely a functional aggregation but a fundamental paradigm shift enabled by large AI models (LAMs), which serve as the ideal engine to break down information barriers and empower intelligent BS agents (IBSAs). Under this framework, perception evolves from raw signal acquisition to semantic understanding, where models like grounded three-dimensional

(3D)-LLMs align heterogeneous sensor data (e.g., lidar, radio-frequency (RF)) for closed-loop reasoning. Communication shifts towards semantic-aware resource management, utilizing these high-level semantic insights, rather than relying solely on channel statistics, to optimize resource allocation and beamforming. Finally, computation is pushed to the edge, where LAMs function as the reasoning core to execute complex planning and autonomous decision-making with low latency. This synergistic relationship, detailed in Table 1, establishes the theoretical foundation for the IBSA architecture.

Traditional small models are typically meticulously customized to solve specific benchmark tasks, relying on carefully tuned hyperparameters and algorithm selection [8]. In contrast, the success of LAMs stems from the scaling effect and the resulting emergent abilities. When model parameters and training data reach a certain threshold, the model’s performance and functionality undergo a qualitative leap, a non-linear growth that traditional small models are unable to achieve [9]. A systematic comparison highlighting these fundamental differences, specifically regarding core capabilities, generalization, and data processing mechanisms, is provided in Table 2.

LAMs pursue general-purpose solving strategies, which is reflected in the research trend toward seeking unified models to address diverse tasks. For example, some research is dedicated to building general-purpose, exemplary deep reinforcement learning (RL) algorithms that achieve competitive performance on a variety of RL benchmarks with a single set of hyperparameters [8]. The generality and zero-shot/few-shot learning capabilities of LAMs enable them to adapt to unstructured, highly variable real-world environments, which is fundamental to achieving general-purpose decision-making and cross-domain collaboration in integrated perception, communication, and computation.


In processing multimodal data, LAMs differ fundamentally from traditional methods. Traditional fusion mechanisms often rely on static feature integration, simply concatenating visual representations from different “expert” encoders, and lack adaptive modulation capability [10]. This rigid combination struggles to effectively address tasks requiring rich multi-scale understanding, such as fine-grained recognition, small object perception, or precise spatial localization. In contrast, multimodal LLM (M-LLMs) can leverage robust cross-modal transfer, in-context learning, and high-dimensional data processing capabilities to achieve dynamic and context-aware heterogeneous data fusion [10]. This mechanism allows the system to flexibly weigh and interpret information from different modalities (such as visual, linguistic, and RF signals) based on current task and environmental demands, thereby enabling more advanced and adaptive decision-making than static integration and laying the technical foundation for real-time, complex decisions in in-

Table 1 Definition of perception, communication, and computation in the LAM-enabled IBSA paradigm.

Function	Core Mechanism via LAM	Strategic Value for 6G
Perception	Semantic Alignment & Grounding: LAMs (e.g., grounded 3D-LLMs) fuse heterogeneous data (lidar, RGB, RF) into unified semantic tokens, translating raw 3D bounding boxes into actionable context.	Enables the BS to “understand” the physical environment rather than just sensing signals, providing the foundation for closed-loop control in tasks like autonomous driving.
Communication	Semantic-Aware Resource Management: The control plane utilizes LAMs to perform context-aware decision-making. Beamforming and resource allocation are driven by high-level semantic understanding of targets (e.g., pedestrians, vehicles) rather than solely by channel statistics.	Transforms the communication system from a signal-processing pipeline into a semantic-aware intelligent agent, significantly enhancing efficiency and reliability in dynamic scenarios.
Computation	Edge Intelligence & Reasoning: Deploying LAMs at the BS edge brings general-purpose cognitive capabilities closer to the user. It supports complex, latency-sensitive tasks such as trajectory prediction, interference coordination, and causal reasoning.	Acts as the “glue” of the integrated perception, communication, and computation loop, enabling the BS to function as an autonomous agent that performs real-time planning and execution without reliance on distant cloud resources.

Table 2 Comparison of LAMs and traditional small models for the integrated perception, communication, and computation paradigm.

Feature	Large AI Models	Traditional Small Models	Strategic Value
Core Capability	General-purpose reasoning, high-level planning, multi-modal semantic alignment.	Domain-specific, feature extraction, optimization for specific tasks.	A unified cognitive and decision-making core.
Generalization	Excellent zero-shot/few-shot capabilities, adaptable to unstructured environments.	Heavily reliant on training sets, poor generalization.	Adapts to the dynamic, highly variable, and non-stationary real-world environment.
Data Processing	Dynamic, adaptive cross-modal fusion mechanisms.	Static feature concatenation, lacks adaptive modulation.	Enables real-time fusion and comprehension of high-dimensional, heterogeneous sensory data.

**Figure 1** The physical evolution of the BS from 1G to 6G and beyond. The architecture has progressed from (Stage 1-2) centralized cabinets (analog BS, digital BTS) with high-loss coaxial feeders to (Stage 3) a distributed structure with a baseband processing unit (BBU) and remote radio unit (RRU) connected by optical fiber. This was followed by (Stage 4-5) the integration of RF and antenna into an active antenna unit (AAU) to support MIMO and massive MIMO. This integration path provides the foundation for (Stage 6) the 6G ISAC BS and (Stage 7) the LAM-enabled IBSA, which leverages multi-modal sensors like lidar and cameras for embodied intelligence.

tegrated perception, communication, and computation.

3. Structure of the Paper

To begin with, Section 2 traces the evolution of BSs from single-function infrastructure to LAM-enabled intelligent agents and provides a systematic comparison with existing next-generation architectures. Subsequently, Section 3 introduces the overall architecture of embodied IBSA. Next, Section 4 elaborates on two key application scenarios of embodied IBSA, namely cooperative vehicle-road autonomous driving (Section 4.1) and ubiquitous embodied IBSA-enabled low-altitude safety (Section 4.2). Following that, Section 5 discusses the enabling technologies that support the implementation of embodied IBSA. Subsequently, Section 6 outlines the metrics and benchmarks designed to evaluate embodied IBSA's performance. Finally, challenges, future directions, and concluding remarks are presented in Section 7 and Section 8, respectively.

II. Evolution of Base Station: From Single-Functional Communication Infrastructure to Multi-Functional LAM-Enabled IBSA

The transition of BS from traditional communication infrastructure to the LAM-enabled IBSA represents a fundamental paradigm shift. This transformation is not merely functional but is mirrored by a profound physical evolution of the BS hardware itself. Over decades, BSs have evolved from centralized analog cabinets with passive antennas to distributed, integrated, and multi-modal intelligent agents. Fig. 1 illustrates this physical evolutionary path, which provides the hardware foundation for the LAM-enabled IBSA that integrates perception, communication, and computation in the 6G era and is central to this paper.

1. Large Model in Communication Network

LAMs are pivotal in driving the evolution of 6G networks due to their versatility and generalization abilities across diverse tasks. Their capacity for cross-task generalization means that these models can be applied to multiple network scenarios without the need for extensive retraining, making them invaluable in dynamic environments [11]. Furthermore, their proficiency in few-shot learning allows for efficient adaptation to new tasks with minimal data, reducing the need for large-scale training datasets [12]. Additionally, LAMs exhibit emergent abilities, where they can develop novel problem-solving strategies that were not explicitly part of their training data [13].

The intersection of large models and edge learning-based agentic models holds immense potential for revolutionizing communication networks. By deploying these models on edge devices, telecommunications systems can enhance both the intelligence and efficiency of resource allocation and management. Research on AI-native telecommunications networks reveals how large models, particularly foundation

models, can empower edge devices to handle complex tasks autonomously [11]. This enables real-time decision-making, reducing latency and improving network reliability. Notably, agentic AI models integrated within telecommunications systems allow BSs to act as intelligent agents that learn from interactions with the environment, leading to self-optimizing networks capable of anticipating and responding to various network conditions [12].

2. Large Model Enabled Embodied Intelligent Base Station Agent

Embodied intelligence refers to systems where intelligence is not just about processing information but is tightly coupled with physical action, perception, and learning within an environment. Traditionally, embodied intelligence has been applied to robotics, where it involves the dynamic interaction of sensors, actuators, and perception to enable real-time adaptation to physical environments [32, 33]. The interaction between the agent and its surroundings forms a continuous loop of sensory input, decision-making, and motor output, providing agents with the ability to learn and adapt [34].

However, this paradigm is now extending beyond robotics and into the realm of communication infrastructure. As networks become increasingly intelligent, embodied intelligence concepts are being integrated into network systems, particularly in 6G environments. This shift allows BSs to act as "embodied agents" that not only process data but actively sense and adjust to the physical network environment, improving network performance through adaptive actions like beamforming, resource allocation, and optimization of spectrum [33]. This transformation is akin to how autonomous vehicles use multi-modal sensor fusion to navigate their surroundings [32]. In a similar manner, network infrastructure can leverage multi-modal sensory data (e.g., radio-frequency perception, visual data) to improve decision-making in dynamic wireless environments [34]. Ultimately, the convergence of embodied intelligence and communication networks has the potential to create more adaptive, resilient systems that are capable of real-time optimization and task-specific decision-making. Such integration marks a fundamental shift in how networks operate, moving from passive infrastructure to intelligent and autonomous systems that can physically interact with their environment [33].

Treating BSs as embodied intelligent BS agents (IBSAs) is an innovative approach in 6G networks. This perspective transforms BSs from mere communication relays to perceptive entities capable of active perception and adapting to their physical environments. With the integration of sensors and AI, BSs can optimize network performance, such as adjusting beamforming based on real-time conditions. However, this shift presents challenges, including the need for seamless multi-modal data fusion and real-time decision-making capabilities that balance computational resources with the network's dynamic demands. Building upon the vision of BSs

Table 3 Representative multi-modality dataset in communication domain. Δ denotes uncertainty.

Ref.	Platform						Sensors/Modality					Weather/Time				Multi-Scenario	Customized Definition	Aerial Target
	CARLA	Sionna	Wireless InSite	Unreal Engine	Real World	Others	Camera	lidar	Radar	CSI	IR	Rain	Snow	Fog	Night			
[14]	×	×	×	×	✓	×	✓	✓	✓	✓	×	✓	×	×	✓	✓	×	✓
[15]	✓	×	✓	×	×	×	✓	✓	✓	✓	×	×	×	×	×	✓	×	×
[16]	×	×	×	✓	×	×	✓	✓	✓	✓	×	Δ	Δ	Δ	Δ	✓	✓	✓
[17]	✓	✓	×	×	×	×	✓	✓	×	✓	×	Δ	Δ	Δ	Δ	✓	✓	×
[18]	×	×	✓	×	×	AirSim, WaveFarer	✓	✓	✓	✓	×	✓	✓	✓	✓	×	✓	×
[19]	×	×	✓	×	×	AirSim, WaveFarer	✓	✓	✓	✓	×	✓	✓	×	✓	×	✓	×
[20]	×	×	✓	×	✓	Blender, Blensor	✓	✓	×	×	×	×	×	×	×	✓	×	×
[21]	×	×	×	×	✓	×	×	✓	×	×	×	×	×	×	×	✓	×	×
[22]	×	×	✓	×	✓	Blender, Blensor	✓	✓	×	×	×	×	×	×	×	✓	✓	×
[23]	×	×	×	×	✓	×	✓	✓	✓	✓	✓	×	×	×	×	✓	×	×
[24]	×	×	×	×	✓	Blender, OpenGL	✓	✓	×	×	✓	×	×	×	×	✓	✓	×
[25]	×	×	×	×	✓	×	✓	✓	✓	×	✓	✓	✓	✓	✓	×	×	×
[26]	✓	✓	×	×	×	Blender	✓	✓	✓	✓	×	✓	✓	✓	×	✓	×	×

Table 4 Representative studies of multimodal perception fusion for communication systems.

Ref.	Platform						Sensors/Modality					Application			
	CARLA	Sionna	Wireless InSite	Unreal Engine	Real World	Others	Camera	lidar	Radar	CSI	IR	Detection, Localization	Beam Prediction	Antenna Position Optimization	Channel Prediction
[16]	×	×	×	✓	×	×	✓	✓	✓	✓	×	✓	×	×	×
[17]	✓	✓	×	×	×	×	✓	✓	×	✓	×	✓	✓	✓	×
[19]	×	×	✓	×	×	AirSim, WaveFarer	✓	✓	✓	✓	×	×	×	×	✓
[20]	×	×	✓	×	✓	Blender, Blensor	✓	✓	×	×	×	×	✓	×	×
[21]	×	×	×	×	✓	×	×	✓	×	×	×	×	✓	×	×
[22]	×	×	✓	×	✓	Blender, Blensor	✓	✓	×	×	×	×	✓	×	×
[23]	×	×	×	×	✓	×	✓	✓	✓	✓	✓	✓	✓	×	×
[24]	×	×	×	×	✓	Blender, OpenGL	✓	✓	×	×	✓	✓	×	×	×
[27]	×	×	×	×	✓	×	✓	×	×	×	×	×	✓	×	×
[28]	✓	×	×	×	×	MATLAB	✓	✓	✓	×	×	×	✓	×	×
[29]	✓	×	✓	×	×	×	✓	✓	✓	✓	×	×	✓	×	×
[30]	✓	✓	×	×	×	×	✓	✓	×	✓	×	×	×	×	✓
[31]	×	×	✓	×	×	×	✓	✓	×	×	×	×	✓	×	×

as embodied IBSAs, there exists two compelling application scenarios: autonomous driving and low-altitude safety. In cooperative vehicle-infrastructure autonomous driving, embodied IBSAs with integrated perception capabilities can play a pivotal role in supporting vehicle-to-everything (V2X) communications through multi-modal data fusion. This allows for real-time, context-aware decision-making that enhances vehicle navigation in complex environments. Similarly, in low-altitude airspace, embodied IBSAs can actively detect and

neutralize potential uncrewed aerial vehicle (UAV) threats. As we move towards 6G, these intelligent, adaptive systems will be critical for ensuring both the safety and efficiency of future networks.

3. Embodied Intelligent Base Station Agent Enabled Autonomous Driving and Low-Altitude Safety

Data is of paramount importance for AI research. However, obtaining real-world data and annotations for cooperative autonomous driving and low-altitude safety scenar-

ios is often exceedingly challenging [25]. Thanks to advances in autonomous driving research, a number of mature simulation environments now support algorithm development and validation within digital twins, including Unreal Engine (a cross-platform, highly extensible real-time 3D engine) [35], CARLA (an open-source autonomous driving simulator) [36], Blender (a free, open-source, cross-platform 3D creation suite) [37], Blensor (a free, open-source simulation package for lidar and Kinect sensors) [38], AirSim (a simulator for drones, cars, and more, built on Unreal Engine) [39], and OpenGL (an open, cross-platform, cross-language graphics rendering application programming interface (API)) [40].

Similarly, the wireless-signal-processing community benefits from its own set of digital-twin tools that lay the groundwork for embodied IBSA: Wireless Insite (a predictive tool for understanding wireless coverage, channel multipath, and data throughput) [41], Sionna (a hardware-accelerated, differentiable open-source library for communications research) [42], and WaveFarer (a high-fidelity radar simulator) [43].

Meanwhile, autonomous driving has produced numerous publicly available multimodal datasets [44, 45]. Recent surveys, such as [46] on lidar and camera fusion, and [45] on radar and camera fusion, provide comprehensive overviews, while works like [47], [48], and [49] push the state of the art in sensor fusion.

Methods for multimodal fusion in autonomous driving offer valuable insights for wireless signal processing. Traditional communications rely on wireless pilot signals alone to estimate the channel and perform demodulation [50–52], but ISAC has emerged as a core 6G capability in which enhanced perception and communications can mutually reinforce one another, enabling applications such as cellular positioning, low-altitude safety, and autonomous driving [53–57]. Table 3 presents an overview of widely used multimodal datasets that serve as the foundation for perception research in the communications domain. Building upon these datasets, Table 4 summarizes representative multimodal perception approaches, including those based on proprietary simulations, real-world measurements, and high-quality public data.

In light of current trends toward large-scale models and embodied intelligence, and building on the trajectory of future communications, the capabilities of existing simulators and datasets, and the wealth of multimodal perception schemes in the literature, this paper explores how embodied IBSA can be most effectively applied in two key areas: vehicle-road cooperative autonomous driving and ubiquitous base-station-enabled low-altitude safety.

4. Comparison with Existing Next-Generation Base Station Architectures

Recent advances at the intersection of LAMs and 6G have led to several architectural proposals that are closely related

to, but still fundamentally different from, the IBSA architecture considered in this paper. Broadly speaking, existing work can be grouped into three lines: (i) LLM/LAM-centric frameworks that endow telecom systems with powerful cognitive capabilities; (ii) network- and edge-centric architectures that distribute large models across a 6G computing continuum; and (iii) agentic or embodied network architectures that explicitly consider interaction with the physical world. This subsection clarifies how IBSA relates to and differs from these lines of research.

The first line focuses on large models as telecom-specific “brains”, but does not explicitly architect a next-generation BS. TelecomGPT is a representative example: it builds a domain-specialized LLM by constructing a telecom corpus, designing pre-training and instruction-tuning pipelines, and defining benchmark tasks such as standard interpretation, troubleshooting, and radio access network configuration assistance [11]. Similarly, the comprehensive survey of LAMs for future communications overviews how LAMs can support a broad spectrum of communication tasks, from channel prediction and beam management to resource allocation and network automation, and sketches a high-level reference framework where LAMs provide AI services to the network [58]. These works clearly articulate the role of LLMs/LAMs as cognitive engines, but the BS is mainly treated as a logical endpoint or data source that queries the model. They do not describe how a single BS should internally organise multimodal perception, LAM-based cognition, and executable actions, nor do they consider safety-critical, closed-loop control at the BS level.

The second line emphasises the network and edge-computing perspective of deploying LAMs. The wireless large AI model framework discusses how a unified LAM can interface with multiple radio functions, proposes an AI-native 6G network architecture, and identifies enabling technologies such as model compression, parallelism, and knowledge distillation [12]. Edge LAM architectures go further by decomposing LAMs and distributing their submodules across geographically separated edge devices, thereby addressing communication, computation, and storage bottlenecks while enabling general-purpose intelligent services [13]. The 6G white paper on LLMs in the computing continuum complements these works by positioning LLMs within a cloud-edge-device hierarchy, where BSs act as part of the edge layer that hosts or accesses LLM-based services and participates in cross-layer orchestration [59]. Taken together, these works provide valuable guidance on where large models should reside, how they can be split, and how they interact with edge resources. However, they largely view the BS as an element of the computing fabric; the internal BS architecture is not the focus, and explicit integration of perception, communication, and local actuation into an embodied BS agent is outside their scope.

The third line explicitly moves from “large models” to “agents” and from passive observation to embodiment. F. Jiang *et al.* in [60] systematically introduced agentic stacks built on top of LAMs, including planning, tool use, memory, and feedback mechanisms, and envisages multi-agent systems for network operation, optimization, and service provisioning. The recently proposed wireless embodied large AI further argues that large models for wireless systems should close the loop with the physical environment via active environmental probing, on-line learning, and interaction with physical interfaces; it outlines design principles and a generic system structure for embodied wireless entities such as user devices or BSs [33]. In parallel, the survey on ISAC towards multifunctional perceptive networks shows how BSs may evolve into perceptive nodes that jointly perform communication and perception by sharing spectrum, hardware, and signal-processing pipelines, and proposes several perceptive-network architectures [61]. These works are highly complementary to IBSA: they highlight the importance of embodiment, agentic behaviour, and integrated perception. Nevertheless, they are either formulated at the level of generic wireless entities or at the network/cluster level, without providing a concrete, BS-centric hardware-software architecture that unifies multimodal perception, LAM-based cognition, and BS-side actuation in specific safety-critical scenarios.

In contrast to the above, the proposed large-model-enabled IBSA architecture is explicitly BS-centric and scenario-driven. Each BS is treated as an embodied agent whose functional focus simultaneously covers communication, perception, and computation. An LAM serves as the unified cognitive core at the BS, while multimodal RF, visual, and point-cloud perception modules feed into this core, and the BS exerts actions through both traditional control-plane commands (e.g., beamforming, power control, scheduling) and environment-shaping wireless interactions (e.g., co-operative perception with vehicles and uncrewed aerial vehicle (UAVs), spectrum-level intervention for low-altitude safety). The architecture is instantiated through a three-layer perception-cognition-execution pipeline, coupled with cloud-edge-end collaboration, digital twins for data generation and simulation, and communication-computation co-design. Unlike purely LAM-centric surveys, IBSA therefore provides a concrete blueprint for how a single BS can be upgraded into an embodied, large-model-enabled agent in specific safety-critical use cases such as cooperative vehicle-road autonomous driving and low-altitude security.

At the same time, IBSA does not aim to replace the above frameworks; rather, it builds on and complements them. Compared with TelecomGPT and LAM surveys [11, 58], IBSA adds explicit multimodal perception and embodiment at the BS level. Compared with edge/continuum LAM architectures [12, 13, 59], IBSA assumes that many of their techniques, e.g., model decomposition, offloading, and life-

cycle management, will be applied under the hood to support the large-model brain of each BS, but these system-level issues are not the main focus of this paper. Compared with WELAI [33], agentic-AI [60], and ISAC-based perceptive-network [61] frameworks, IBSA specialises their general principles to a next-generation BS design that tightly couples integrated perception-communication-computation with an LAM-based cognitive core and concrete safety-critical scenarios. As summarised in Table 5, representative works from these three lines and the proposed IBSA are compared along four dimensions: (i) functional focus of the next-generation BS (communication, perception, computation); (ii) type of BS-side core intelligence (LAM, LLM*, or no unified large model); (iii) whether the BS only passively receives wireless signals or can actively interact with and shape the environment via wireless actions; and (iv) targeted key scenarios (autonomous driving, low-altitude security). Table 5 highlights that existing architectures typically cover only a subset of these dimensions, whereas IBSA jointly realises all four and thus provides a concrete instantiation of a large-model-enabled embodied intelligent BS for integrated perception-communication-computation 6G systems.

III. Architecture of Intelligent Base Station Agent

1. Overall Architecture and Deployment Forms of Embodied IBSA

With the evolution of 6G toward “native intelligence” and “perception-communication-computation integration,” BSs cease to be mere passive forwarding nodes and become embodied IBSAs endowed with perception, cognition, and actuation capabilities. The overall architecture of IBSA follows an upward three-layer closed-loop:

- Perception layer: acquires and pre-processes multimodal data.
- Cognition layer: performs cross-modal fusion, world modeling, and policy generation.
- Execution layer: carries out controllable actions on wireless resources and external entities.

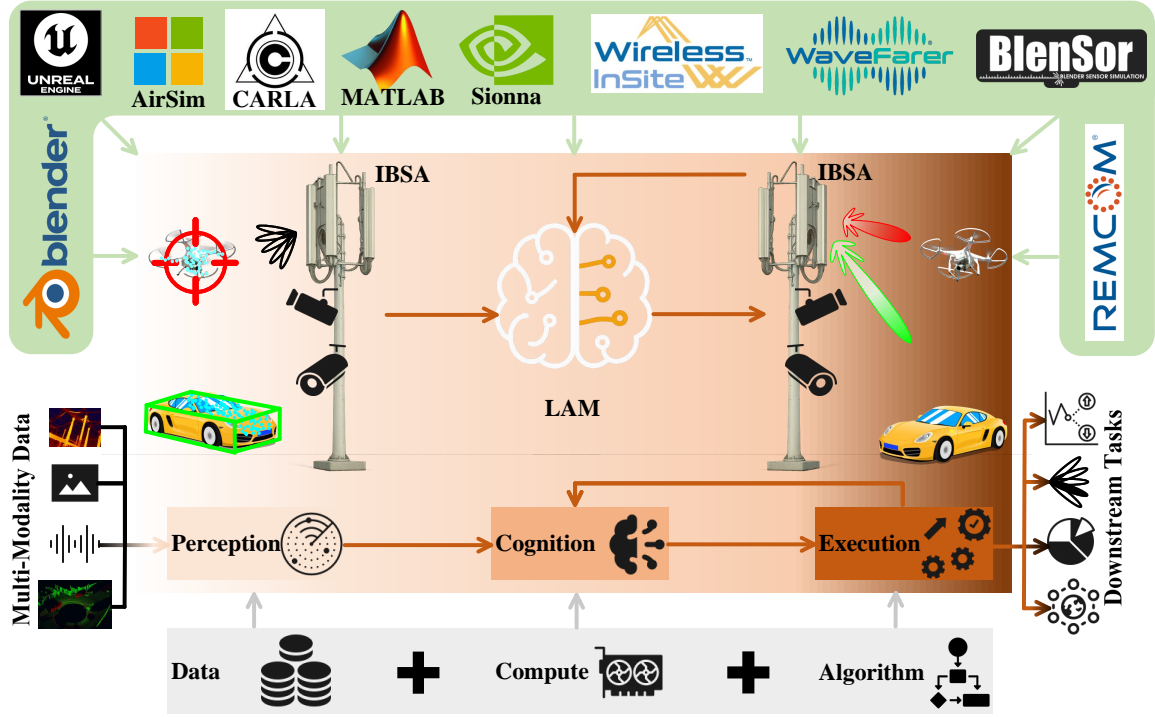
Above these three layers, a LAM provides a unified representation and inference interface, and is deployed across cloud-edge-end in a collaborative fashion to satisfy varying latency and energy constraints. As illustrated in Fig. 2, upper-level multi-engine simulation environments (e.g., Unreal Engine, Blender, AirSim, CARLA, MATLAB, Sionna, Wireless InSite, WaveFarer, Remcom, etc.) continuously feed multimodal data and digital-twin feedback downward; multiple embodied IBSA nodes in the mid-tier complete perception-cognition-execution cycles under the drive of the LAM; and the lower tier’s data, compute resources, and algorithmic foundations underpin the entire agent ecosystem.

The perception layer, visually represented on the left in Fig. 2 as “Multi-Modality Data” flowing into the

*The difference between LAM and LLM is detailed in [62].

Table 5 Systematic comparison between representative next-generation BS-related large-model/ISAC architectures and the proposed large-model-enabled IBSA.

Ref.	Functional Focus of Next-generation BSs			Core Intelligence at the BS		BS-environment Interaction via Wireless Signals		Key Scenarios for Next-generation BSs	
	Comm.	Perception	Computation	LAM	LLM	Passive Reception Only	Active Interaction	Autonomous Driving	Low-altitude Safety
[11]	✓	✗	✗	✗	✓	✓	✗	✗	✗
[12]	✓	✓	✓	✓	✗	✓	✗	✗	✗
[13]	✓	✓	✓	✓	✗	✓	✗	✗	✗
[33]	✓	✓	✓	✓	✗	✗	✓	✓	✗
[58]	✓	✗	✓	✓	✗	✓	✗	✗	✗
[59]	✓	✗	✓	✗	✓	✓	✗	✓	✗
[60]	✓	✗	✓	✓	✗	✓	✗	✗	✗
[61]	✓	✓	✓	✗	✗	✗	✓	✓	✓
Ours	✓	✓	✓	✓	✗	✗	✓	✓	✓

**Figure 2** Architecture of embodied intelligent BS agent. The framework is centered around an LAM that empowers the embodied IBSA through a three-layer architecture, consisting of a perception layer, a cognition layer, and an execution layer. At the top, various simulation platforms are used to synthesize and replay scene data for closed-loop evaluation. The agent processes multi-modality data inputs, such as infrared, visual images, wireless signals, and lidar point clouds. This enables the embodied IBSA to perform various downstream tasks, which include prediction, beam-forming, resource allocation, and network cooperation. The entire architecture is supported by the foundational components of data, compute resources, and algorithms.

“Perception” stage, emphasizes “visible,” “audible,” and “measurable” perception. Inputs from the wireless domain include uplink/downlink pilots and channel state information (CSI), millimeter-wave/terahertz (THz) radar echo matrices, and angle-of-arrival/time-difference-of-arrival (AoA/TDoA) estimates; external sensors comprise roadside and vehicle-mounted cameras, lidar, and thermal infrared im-

agers. Multimodal data must be filtered, compressed, and timestamped at the acquisition point and spatially registered to a common reference coordinate system. To control network load, event-driven uploading, learned compression, or feature-level upload strategies may be adopted. The perception layer also outputs data-quality and uncertainty estimates, furnishing the cognition module with quantified confidence

measures.

The cognition layer, which is fundamentally embodied by the LAM, functions as the “central brain” of embodied IBSA, symbolized by the central LAM brain icon in Fig. 2. At its core lies a multimodal large model with vision-language-radio coupling capability, which maps heterogeneous signals into token sequences and captures intermodal mutual information and constraints via cross-modal attention and spatiotemporal Transformers. Training draws on both real-world network and road/airspace data (e.g., vehicle and UAV control logs) and long-tail scenarios generated by digital-twin platforms, employing self-supervised or contrastive learning to bolster generalization and robustness. Mechanisms such as federated learning, parameter-efficient fine-tuning (PEFT), and mixture-of-experts (MoE) enable model sharding across cloud, edge, and end devices, thus achieving privacy protection and low-latency inference. The cognition layer should further incorporate memory and causal-reasoning capabilities to support long-term state tracking, threat-intent analysis, and policy interpretability.

The execution layer, the final stage in the pipeline shown in Fig. 2, implements cognitive outputs as controllable actions in both the communications and physical domains. On the wireless side, embodied IBSAs may perform beam-forming, power/spectrum scheduling, network-slice allocation, and reconfigurable intelligent surface (RIS) configuration. On the security/control side, they can target unauthorized (“black-fly”) drones or anomalous vehicles with directional jamming, link spoofing, and coordinated behavior control (e.g., formation scheduling, obstacle-avoidance path broadcasting), feeding execution outcomes back to the cognition layer to close the optimization loop [63]. Execution policies must balance multiple objectives, often within reinforcement-learning or game-theoretic frameworks and supplemented by safety constraints and human-machine collaboration interfaces to ensure controllability: communication quality of service (QoS) versus perception accuracy, jamming efficiency versus collateral damage, automated decision-making versus legal compliance.

From an engineering and deployment standpoint, embodied IBSA can be classified by entity type, application domain, and deployment form. Fixed ground BSs are suited for continuous urban and highway coverage, whereas aerial BSs (UAV-BS, balloon-borne) provide rapid blind-spot fill-in during disasters or large events [64–66]. Application domains span V2X cooperative perception, UAV control, smart-city security, and industrial internet-of-thing (IoT), each with distinct latency, reliability, and privacy requirements. Deployment follows a cloud-edge-end continuum: the cloud hosts ultra-large models for retraining and global updates; the edge runs model shards or expert subnetworks for latency-sensitive fusion inference and resource orchestration; end devices execute distilled or pruned lightweight models to enable de-

graded operation when disconnected. These tiers maintain consistency via parameter synchronization, knowledge distillation, and experience replay, while model routing and caching strategies alleviate network pressure.

The embodied IBSA architecture, through its three-layer closed-loop of perception, cognition, and execution, tightly couples the triad of data, compute, and algorithms with the pathway of multimodal data, large models, and downstream tasks. Digital twins serve as data amplifiers for training and safe sandboxes for policy validation; large models provide cross-task shared representations and few-shot adaptation capabilities; and integrated communication/perception hardware with network orchestration ensures efficient real-world realization of cognitive outcomes.

2. Multimodal Large Model at the Cognition Layer

The cognition layer of an embodied IBSA is fundamentally embodied by a multimodal LAM that acts as the “central brain”. Beyond the high-level architecture described in Section 3.1, we provide here a representative instantiation of such a model, consistent with recent multimodal sensing-and-communication datasets [14–16, 18, 19, 25] and multimodal fusion architectures for autonomous driving and V2X [17, 23, 24, 27, 30, 31, 45–49].

1) Model structure: The multimodal LAM follows a Transformer-style encoder-decoder architecture with heterogeneous modality encoders and a shared spatiotemporal backbone. Wireless-domain inputs (uplink/downlink CSI, millimeter-wave (mmWave)/THz radar tensors, AoA/TDoA estimates) are first pre-processed into channel and radar feature tensors following the data formats and synchronization strategies adopted in multimodal sensing-and-communication datasets such as DeepSense 6G [14], Deepverse 6G [15], M3SC [18], SynthSoM [19], and V2X-radar [25]. These tensors are then fed into complex-valued convolutional or recurrent encoders to produce one-dimensional “radio tokens”.

Roadside and vehicle-mounted cameras are encoded by a vision Transformer (ViT)-style backbone into “vision tokens”, while lidar point clouds are processed by voxel-based or point-based encoders as in state-of-the-art sensor fusion frameworks [45–49]. Each token is augmented with learnable modality-type, spatial, and temporal embeddings so that heterogeneous inputs share a unified token space.

The resulting token sequences from all modalities are concatenated and passed through a stack of spatiotemporal Transformer blocks. Within each block, self-attention operates over the full multimodal sequence, while cross-modal attention explicitly enhances interactions between radio, vision, and point-cloud tokens, following the paradigm of MVX-ViT [17] and related fusion architectures such as BEVFusion, MSSF, and CR3DT [47–49]. On top of the shared backbone, multiple task-specific heads are attached for 3D object detection and tracking, link-quality or beam-index prediction,

jamming/beamforming policy generation, and trajectory forecasting, similar in spirit to recent vehicular multi-sensor fusion designs [45–49].

To make the design compatible with practical edge deployment, the IBSA-side LAM is designed to contain on the order of 10^9 trainable parameters, which is comparable to compact multimodal Transformer models used in real-time perception and communications [12]. Larger cloud-resident teacher models in the 10^{11} – 10^{12} parameter range can be used for pre-training and distillation, while edge-side models are obtained via model compression and parameter-efficient fine-tuning [58].

2) Training data sources and learning objectives: The multimodal LAM is trained on a mixture of real and synthetic data. On the real-data side, we leverage public multimodal datasets that provide synchronized camera, lidar, and radar measurements [25, 44, 45], together with communication-oriented datasets such as DeepSense 6G and M3SC that align CSI or radar signals with visual and point-cloud observations [14, 18]. On the synthetic side, joint traffic–wireless digital twins based on tools such as Wireless InSite, Sionna, and WaveFarer [15, 41–43] are used to generate long-tail safety-critical scenarios and rare channel conditions that are difficult to capture in the field.

Training combines self-supervised objectives (e.g., masked token modeling and cross-modal contrastive learning) with supervised detection, localization, and beam-prediction losses, following recent multimodal learning practices in autonomous driving and ISAC [45–49]. Federated learning, parameter-efficient fine-tuning, and MoE strategies [60] further enable model partitioning across cloud–edge–end, providing privacy preservation and low-latency inference.

3) Integration of heterogeneous modalities: To integrate heterogeneous modalities in a principled manner, the LAM first performs spatiotemporal alignment at the feature level: raw CSI, images, and point clouds are time-stamped and projected into a common world coordinate frame, as recommended in recent surveys on multimodal fusion [45–49]. After tokenization, the model interleaves radio, vision, and point-cloud tokens within the same sequence and lets the Transformer attend over the full set. Cross-modal attention allows, for example, occluded vehicles that are only visible in radio-domain CSI to be associated with partially visible objects in the image space, while temporal attention stabilizes decisions across successive frames. In this way, the cognition layer naturally fuses “through-the-obstacle” RF cues with high-semantic visual and geometric information, and propagates the fused representation to downstream communication and control heads.

IV. Key Application Scenarios

Building on the general framework outlined above, we now turn to two illustrative application scenarios. The first demonstrates how multi-source fusion perception empowers robust, all-weather autonomous driving in complex traffic environments. The second shows how ubiquitous embodied IBAs deliver full-band, all-modal safety monitoring and response in low-altitude UAV operations.

1. Cooperative Vehicle-Road Multi-Source Fusion Perception Empowering Autonomous Driving

In urban canyons, highway on-ramps, tunnel entrances, and traffic environments dense with large vehicles, single-modality perception exhibits structural limitations: millimeter-wave/THz links, owing to their strong directivity, are highly sensitive to blockage (when vehicles, pedestrians, or buildings intervene, link quality and detection capability degrade sharply) [67]; cameras, although offering high spatial resolution and rich semantic content, are acutely vulnerable to illumination changes and adverse weather, with nighttime glare, backlighting, and rain- or fog-induced scattering frequently causing missed or false detections; lidar and conventional radar are respectively constrained by cost and angular resolution. Hence, the development of multi-source information fusion mechanisms via vehicle-road-base-station cooperation is widely recognized as the key to achieving all-weather, all-element safety in autonomous driving [17].

Against this backdrop, the complementarity between cellular BSs and vision/point-cloud sensors becomes increasingly pronounced. BSs, operating in the sub-6 GHz, upper-6 GHz centimeter-wave, millimeter-wave, and even THz bands, benefit from stable power supplies, wide-area coverage, and deployable edge computing resources; by exploiting physical-layer features such as channel state information (CSI) and beam-training feedback signals, they can achieve non-line-of-sight “through-the-obstacle” perception. Roadside and vehicle-mounted cameras and lidar, in turn, furnish high semantic density and fine spatial detail over local regions. These modalities complement one another in spatial scale (wide-area vs. local), information dimension (physical channel vs. visual semantics), and robustness (occlusion vs. illumination): the BS can provide a priori target regions and velocity estimates within visual blind spots, while the visual modality feeds back to the communication side to refine beam selection, interference management, and resource scheduling, thereby realizing a co-optimization of “visibility” and “connectivity” [17, 29].

To unlock this complementarity, the LAM’s inherent capabilities for cross-modal alignment and shared representation learning are central. Models such as MVX-ViT (Multimodal/V2X vision Transformer) in [17] encode wireless features into “radio tokens” and image or point-cloud features into “vision tokens,” then employ cross-modal atten-




Figure 3 Cooperative vehicle-road multi-source fusion perception empowering autonomous driving. This figure illustrates the complete closed-loop workflow of multi-source fusion perception across the scenario, data, and network layers. The entities on the left, “Vehicle, IBSA, Camera, and Cloud Server”, delineate the primary roles in the system. Within the road scenario, the colored arrows represent data flows: blue arrows denote the uplink of environmental information from the embodied IBSA to the cloud; red arrows signify the acquisition and transmission of visual data from cameras; and green arrows indicate robust beamforming and high-fidelity communication from the embodied IBSA to vehicles, guided by perception results and cloud directives. The workflow on the right begins with the collection of “Multi-Modality Data”, including RF signals, visual images, and lidar point clouds. This data then proceeds through a vertical pipeline of “Spatiotemporal Alignment”, “Edge Inference”, and “Network Cooperation” to achieve perception-communication-computation co-optimization for autonomous driving tasks. The overall diagram highlights the embodied IBSA’s pivotal role as a perception-communication bridge within this closed-loop system.

tion across spatial and temporal dimensions to establish robust associations even under occlusion or adverse illumination [17]. These architectures typically include multi-task output heads that jointly predict object detection/tracking results and key communication metrics (e.g., optimal beam indices, link quality), thereby embedding perception feedback directly into network resource scheduling. Furthermore, by leveraging vision-radio consistency and temporal coherence constraints, self-supervised or contrastive learning strategies enable cross-modal alignment with minimal large-scale annotations, substantially reducing labeling costs while enhancing generalization [27].

As illustrated in Fig. 3, the entire multi-source fusion workflow forms a coherent closed-loop across the scenario layer, data layer, and compute/network layer. On the left, four abstracted entities, vehicle, embodied IBSA, camera, and cloud server, delineate the roles of the in-vehicle unit, embodied IBSA, roadside camera, and cloud/edge server. In a real-road scenario, blue arrows denote the uplink of environment information perceived by the embodied IBSA; green arrows indicate robust beamforming and high-fidelity data communication performed by the embodied IBSA based on perceived vehicle targets and cloud directives; red arrows mark the acquisition and transmission of visual image streams. “Multi-Modality Data” comprises visual or infrared images from roadside cameras, wireless signals received by the embodied IBSA, and lidar point clouds. The system

then advances vertically through “Spatiotemporal Alignment → Edge Inference → Network Cooperation”: first realizing cross-device and cross-modal alignment in space and time; next executing fusion inference at the edge (e.g., via a lightweight MVX-ViT or MoE expert network); and finally distributing results through V2X interfaces to achieve perception-communication-computation co-optimization for autonomous driving tasks. The color coding and hierarchical framework visually convey modality complementarity, spatiotemporal synchronization, and the causal linkage between edge intelligence and network cooperation, while underscoring the embodied IBSA’s pivotal role as a perception-communication bridge within the closed-loop system.

To validate the effectiveness of this architecture, performance evaluation must span perception, communication, and agent dimensions. Perception metrics include mean average precision (mAP), recall rate, distance and velocity estimation errors, and cross-modal matching accuracy; communication metrics encompass end-to-end latency (covering acquisition, synchronization, inference, and result dissemination), top- k beam selection hit rate, link throughput and reliability, as well as uplink perception data volume and downlink broadcast frequency (network load); from the agent’s perspective, key indicators include task completion rate, collision-rate reduction, model parameter count and energy consumption, and interpretability. A unified key performance indicator (KPI) framework is essential to quantify overall benefits and to

guide decisions on model compression, compute scheduling, and bandwidth allocation [31].

Existing research can be categorized into three main threads. First, multimodal cooperative perception datasets and benchmarks (e.g., MVX-Dataset, MultiX [68]) provide annotation and evaluation frameworks for joint vision-radio tasks but predominantly focus on vehicle-to-infrastructure (V2I) scenarios and lack large-scale public base-station-level wireless raw data [69]. Second, digital-twin-driven simulation platforms are maturing: by coupling CARLA with wireless ray-tracing tools such as Sionna or Wireless InSite, they enable real-to-twin close-loop optimization; representative studies [29, 31] explore beam selection and latency-accuracy trade-offs. Third, large-model architectures, including MVX-ViT and radio-vision Transformer, demonstrate significant potential but remain challenged by parameter scale, inference latency, and edge-deployment constraints, leaving large-scale real-world deployment still out of reach [17]. Overall, unified interfaces, standardized evaluation protocols, and real-road validation remain underdeveloped in the public domain.

Despite these advances, multiple challenges persist in data privacy and security, standardized interfaces, and real-world deployment trials. Vehicle and roadside video often contain sensitive personal information, necessitating advanced privacy-preserving mechanisms far beyond basic data masking. While federated learning and differential privacy provide a baseline by keeping data local and adding noise, they may not sufficiently protect against inference attacks on shared model gradients. To address this, robust anonymization schemes are required. Homomorphic encryption is a powerful tool, allowing the IBSA or cloud server to perform computations (e.g., model aggregation) directly on encrypted data without decryption, thus ensuring sensitive trajectories and features remain confidential even from the coordinator [70]. Further enhancements include integrating federated learning with secure multi-party computation to distribute the aggregation task, or using zero-knowledge proofs to allow vehicles to prove they possess valid data without revealing the data itself [71].

Wireless data likewise may reveal trajectories and identities, requiring a robust combination of these cryptographic, policy, and technical safeguards[27]. Concurrently, the absence of industry standards for cross-modal data formats, synchronization protocols, and model APIs impedes multi-vendor collaboration and large-scale adoption. Transitioning from laboratory to roadway demands system-level trade-offs among fifth-generation (5G)/6G network slicing, edge compute scheduling, device heterogeneity, and cost, alongside the establishment of open testbeds and sandbox policies [31]. Furthermore, to address dynamic variations in construction, seasons, and traffic flow, models must support continual learning and domain adaptation; under con-

strained bandwidth and compute resources, joint optimization of sampling rates, upload strategies, inference accuracy, and scheduling policies will be the core challenge for the long-term evolution of multi-source fusion-enabled autonomous driving [27, 29].

2. Ubiquitous Intelligent Base Stations Agent Empower Low-Altitude Safety

Low-altitude airspace is rapidly transitioning from sparse control to dense openness, and the widespread use of UAVs in logistics, inspection, and emergency response brings real threats of unauthorized flights and airspace incursions. A typical adversarial scenario involves small, uncrewed rotorcraft loitering over urban cores, energy hubs, or large-scale events; their low radar cross-section, agile flight paths, and sudden appearances evade traditional single-modality monitoring. Any misclassification or missed detection incurs substantial legal liability and social cost. Consequently, anchored by embodied IBSA with their wide-area coverage and edge computing advantages, constructing an all-weather, full-band, multi-modality low-altitude security perception and response closed-loop has become a critical component of the “new air traffic management” system.

At the heart of this objective lies the fusion of BS side wireless perception with heterogeneous modalities such as vision and acoustics. On the wireless side, uplink/downlink pilots, CSI anomaly patterns, millimeter-wave/THz radar echoes, and AoA estimates enable non-line-of-sight detection of high-speed or occluded targets. Vision and thermal-infrared cameras supply fine-grained shape information and rotor-blade spectral cues, while microphone arrays capture airframe-specific acoustic signatures. Multi-modality collaboration maintains high detection confidence even when a single modality fails (e.g., at night, in dense fog, or under strong interference), and mutual corroboration suppresses false alarms. Datasets such as DeepSense 6G in [14] demonstrate the importance of synchronized acquisition and temporal alignment for multi-modal communications-perception paradigms, offering reusable templates for data-structure design and unified evaluation metrics; Great-X/Great-MSD single-engine simulator in [16] realizes synchronous CSI, RGB, and lidar generation in low-altitude UAV scenarios, validating CSI-driven 3D localization and cross-engine generalization, and providing high-fidelity support for “real-twin” co-modeling and algorithm verification.

As depicted in Fig. 4, this subsection’s schematic illustrates the embodied IBSA-enabled low-altitude security closed-loop across three dimensions: entity roles, information flows, and decision chains. At the scene center, green arrows denote continuous monitoring and data uplink for authorized UAVs, while red arrows indicate early warning and targeting of suspicious or intruding UAVs. Orange and yellow vertical arrows mark data uploads from embodied IBSA/cameras to edge servers, and white bidirectional arrows show collabora-




Figure 4 Ubiquitous embodied IBSA empower low-altitude safety. This figure illustrates the closed-loop process for low-altitude security enabled by the IBSA, detailing entity roles, information flows, and the decision chain. In the central scenario, green arrows represent the continuous monitoring and data uplink for legitimate UAVs, while red arrows signify the early warning, targeting, and eventual jamming of unauthorized or suspicious UAVs entering a restricted area. The vertical arrows indicate data uploads from embodied IBSAs and cameras to the edge/cloud server, and the white bidirectional arrows depict the crucial role of “Network Cooperation” among multiple embodied IBSAs and cloud-edge nodes for tasks like collaborative localization and resource scheduling. The workflow on the right outlines the full decision chain, progressing from “Detection and Monitoring”, through “Evaluation and Decision”, to responsive actions like “Jamming” or issuing control commands to neutralize threats. This schematic underscores the embodied IBSA’s central function in an integrated system that provides perception, assessment, and interference capabilities for low-altitude airspace management.

tion among multiple base-stations and cloud-edge nodes. On the right, the “Detection and Monitoring → Evaluation and Decision → Jamming” pipeline captures the full chain from perception through assessment to interference or hijacking. The diagram underscores the pivotal role of the “Network Cooperation” module in multi-point collaborative localization, interference resource scheduling, and strategy dissemination [72].

At the algorithmic level, LAMs furnish unified representations for UAV identification, trajectory prediction, and threat assessment [73]. Heterogeneous inputs, such as wireless CSI sequences, visual frames, acoustic spectrograms, are tokenized into a common sequence, and cross-modal attention with temporal modeling performs identity classification and intent inference. Generative or world models (e.g., Great-X/Great-MSD, DeepSense’s digital-twin scenarios) can synthesize “long-tail threat trajectories” and “complex interference environments” within simulation domains, enabling risk-adaptive evaluation and game-theoretic strategy optimization. To meet continual learning demands, federated learning, PEFT, and incremental contrastive learning support online model updates and domain generalization without uploading raw sensitive data.

Designing countermeasures requires balancing effectiveness, controllability, and compliance. For jamming, reconfigurable intelligent surfaces (RIS) or coordinated multi-BS beamforming can apply directional power suppression or

link spoofing to a target UAV, avoiding indiscriminate electromagnetic emissions and collateral damage. For hijacking or forced landing, precise localization confirmed by vision/acoustics enables low-power global position system (GPS) or command-link spoofing to guide the UAV into a safe landing zone. Localization and tracking leverage multi-BS TDoA and AoA fusion with particle filtering to reconstruct trajectories at centimeter-level accuracy, supporting forensic evidence and post-operation analysis. It must be stressed that such active interference is not a standard network function. Under international frameworks such as the international telecommunication union (ITU) radio regulations [74], intentional jamming is broadly prohibited to prevent harmful interference with legitimate radiocommunication services. It is imperative that any active interference therefore operates as a specific, national-level legal exception, strictly under emergency authorization (e.g., by public safety or defense authorities) and within designated frequency bands, as discussed below.

Regarding experimental benchmarks and public datasets, real low-altitude multi-modal data remain scarce. To date, only Great-MSD in a single-engine environment has constructed large-scale synchronized UAV trajectory, CSI, and vision datasets, providing benchmark experiments for CSI-driven 3D localization. DeepSense 6G offers limited UAV scenarios but primarily focuses on V2I communications and ground-based perception. There is an urgent need

for standardized low-altitude security data formats (including threat labels, interference strategies, and legal statuses), unified evaluation metrics (detection rate, false alarm rate, and response latency triplets), and cross-domain transfer and adversarial-robustness assessment protocols.

Although multi-station architectures exhibit significant potential, they face multiple challenges. First, legal compliance and liability allocation: UAV detection, jamming, and forced landing involve air-traffic control, public security, and radio-management authorities. As noted, the act of jamming is generally illegal under ITU regulations [74]. Therefore, the “authorization” process mentioned is a critical legal procedure, typically reserved for designated national agencies (e.g., ministry of defense, national police) to counter immediate security threats. This capability cannot be deployed by civil or commercial IBSA operators without such explicit, lawful delegation [75]. Balancing this emergency response authority with privacy protection demands a technology-policy co-design. Second, false-alarm costs and trust calibration: misclassifying a legitimate UAV as a threat entails economic and legal risks, necessitating confidence thresholds, self-auditing, and human-in-the-loop review mechanisms. Third, collaborative asynchronous transfer mode and standard interfaces: the lack of cross-operator and cross-vendor data interfaces and synchronization protocols limits multi-BS cooperation and cross-domain coordination. Fourth, real-world deployment trials: edge-compute overhead, network slicing orchestration, device heterogeneity, and cost control require phased validation from simulation through testbeds to live urban environments. Looking ahead, establishing open low-altitude safety sandboxes, advancing privacy-preserving multimodal data computation, and developing joint communication-perception-control resource orchestration algorithms will be the key pathways for “ubiquitous BSs” to ensure low-altitude security.

3. Generalizability Analysis of Heterogeneous IoT Scenarios

While the preceding subsections focus on autonomous driving and low-altitude safety, the proposed LAM-enabled IBSA architecture possesses inherent generalization capabilities that extend to other critical 6G scenarios, most notably industrial IoT and smart cities. The core “perception-cognition-execution” closed-loop, driven by the multimodal fusion capabilities of LAMs, serves as a universal solution for heterogeneous environments characterized by visual occlusion and complex dynamics.

1) Industrial IoT

The smart factory represents an extreme environment that challenges traditional perception systems. Unlike open roads, factories are dense with metallic racks, moving robotic arms, and high-density equipment, creating severe visual occlusion and complex multipath wireless propagation. Autonomous mobile robots (AMRs) often face blind spots that onboard

perception devices (e.g., lidar) cannot resolve. In this context, the IBSA functions as critical infrastructure for the factory, generalizing the “vehicle-road cooperation” model (Section 4.1) into “AMR-infrastructure collaboration”. Leveraging the “through-the-obstacle” nature of wireless signals, the IBSA fuses data from surveillance cameras with uplink CSI and millimeter-wave radar perception. This allows the system to perceive AMRs or humans hidden behind metallic obstructions, where visual perception fails. Furthermore, the edge-deployed LAM processes this multimodal stream to construct a real-time, dynamic “world model” of the factory floor, predicting the trajectories of both humans and machines. Replicating the safety loop of low-altitude scenarios (Section 4.2), the IBSA ensures human-machine safety via its execution layer. Upon predicting a potential collision, the system does not merely alert but triggers millisecond-level commands (e.g., emergency braking for robotic arms) via ultra-reliable low-latency communication, realizing a true integration of perception, communication, and computation.

2) Smart City

At a macroscopic level, the IBSA architecture scales to support smart city applications. Here, the IBSA acts as a ubiquitous intelligent node, an “urban neuron”—managing diverse data streams ranging from traffic flows to smart grid loads. The traffic management optimization (e.g., “Green Wave” for emergency vehicles) or intelligent load balancing strategies in smart grids follows the same “perception-cognition-execution” paradigm found in mobility scenarios. The LAM at the edge serves as a regional brain, fusing heterogeneous data (visual traffic data, power consumption metrics) to generate optimal control policies for resource allocation.

In summary, the applicability of the IBSA is not limited to specific verticals. Its ability to convert multimodal perception data into intelligent decisions and close the loop through wireless execution proves its potential as a generalized infrastructure for the heterogeneous scenarios of the 6G era.

V. Key Enabling Technologies of Embodied IBSA

In the previous section, we established the overall architecture of embodied IBSA, clarifying its large-model-centric design, the closed “perception-cognition-execution” loop, and its reliance on cloud-edge-end collaboration. However, to realize this architecture in practice and ensure stable operation in complex, dynamic environments, a suite of key enabling technologies must work in concert across computation, data, algorithms, security, and governance.

Among these, the LAM itself forms the cognitive heart of the IBSA. To detail this core component, Table 6 provides a specific summary of the key machine learning techniques that empower this “central brain,” from its foundational architecture to its training and inference methods [9, 10, 12, 13, 33, 62, 76].

Complementing this, Table 7 provides a high-level

overview of the key system-level technologies discussed in this section, outlining their respective roles from perception and simulation to resource orchestration and security. The following subsections will now elaborate on these critical topics in detail.

1. Cognitive Core: Algorithms for Perception and Reasoning

1) Advanced Perception and Fusion Capabilities

To achieve robust environmental perception, as required by the first layer of the architecture in Fig. 2, the embodied IBSA overcomes the limitations of individual sensors. It fuses high-resolution data from external cameras and lidar with its own unique ability to sense using communication signals, which can penetrate obstacles and adverse weather that impair visual systems. This fusion process represents the critical first stage of the perception–cognition–execution optimization loop, where the choice of data modality and fusion complexity directly impacts the subsequent computational load (cognition) and the timeliness of responsive actions (execution). To make this fused data actionable for safety-critical tasks, all computation is performed at the network edge, enabling millisecond-level inference and decision-making. This edge deployment is essential for managing the inherent trade-offs between perception accuracy and latency within the closed-loop. Moreover, the system’s reliability in unpredictable, “long-tail” scenarios is continuously hardened by training on vast datasets generated within digital twin environments and by using self-supervised learning to enhance generalization.

2) Edge LAM Training and Inference Optimization

Embodied IBSA’s large-model core must perform low-latency inference and on-demand fine-tuning at the edge. Parameter-efficient fine-tuning methods such as low-rank adaptation (LoRA) train only low-rank adapter matrices while keeping the backbone weights frozen, enabling rapid domain adaptation and significantly reducing memory footprint and communication overhead. The MoE architecture sparsely activates expert subnetworks, mapping different wireless scenarios and task types (perception/communication/control) to specific experts for “compute-on-demand” allocation. To prevent routing oscillation and load imbalance, temperature-controlled top- k gating and load-balancing losses can be employed [77]. Hierarchical scheduling addresses heterogeneous cloud-edge-end resources. Scheduling policies should jointly consider graphics processing unit (GPU) utilization, link conditions, and task urgency, and may leverage reinforcement learning or game-theoretic optimization for bandwidth-computation co-orchestration.

3) Interpretability Considerations

Since embodied IBSA’s decisions directly impact traffic safety and low-altitude management, interpretability is a prerequisite rather than an afterthought. Beyond stan-

dard attention visualization and feature-attribution methods (SHAP [78], Integrated Gradients [79]), structured explanations for spatiotemporal sequences and multimodal fusion are required: explicit wireless-vision alignment weights, expert selection paths, and beamforming adjustment causal chains. Furthermore, symbolic reasoning or causal graph models can produce verifiable intermediate conclusions, and a natural-language explanation head can articulate decision rationales to operators, enhancing human-machine collaboration and regulatory transparency.

2. System Architecture: Coordination and Orchestration

1) Cloud-Edge-End Coordination and Model Compression Strategies

To accommodate heterogeneous computational and storage capabilities across nodes, model compression techniques, such as pruning, quantization, knowledge distillation, and tensor decomposition, must be combined. End devices retain only task-specific heads or lightweight transformers, edge nodes maintain medium-scale backbones with online update capability, and the cloud preserves full-parameter models and version histories. Coordination mechanisms must address consistency (version conflicts, cache staleness), hardware heterogeneity, and link instability. Gradient compression (e.g., top- k [80], random sparsification [81]) and joint encoding-scheduling design ensure timely synchronization under limited uplink bandwidth [82].

2) Resource Orchestration: Network Slicing and Perception-Communication-Computation Co-optimization

Embodied IBSA concurrently supports communication services, perception tasks, and intelligent computation, which forms a complex perception–cognition–execution collaborative optimization problem. This requires network slicing to partition spectrum, time slots, cache, and GPU cores into logical resource pools for dynamic task orchestration under diverse QoS/quality of experience requirements. This co-optimization can be formulated as a joint-constraint multi-objective problem: minimize end-to-end latency and energy consumption while maximizing perception accuracy, cognitive inference quality, and communication throughput [83]. Solutions may leverage distributed Lagrangian decomposition, deep reinforcement learning, or graph neural networks on large-scale topologies to manage these complex trade-offs [12, 58]. For bursty tasks (e.g., UAV intrusion), preemptive slice reconfiguration and priority scheduling must be supported to guarantee real-time performance for critical missions.

3. Security and Trustworthiness: Privacy and Defense

1) Federated Learning and Privacy Protection

When multiple BSs, vehicles, and cameras collaboratively update the distributed LAMs, raw data are often constrained by privacy and bandwidth. federated learning enables nodes to upload only gradients or model deltas; combining adap-

Table 6 Key LAM technologies for IBSA.

Category	Key Technology	Introduction & Purpose
Infrastructure	Transformer	Cornerstone of LLMs. Uses self-attention for parallel processing and long-range dependency capture. Purpose: Processing 6G time-series signals (e.g., CSI, radar waveforms) and modeling complex spatiotemporal correlations.
	SSM (State Space Model) / Mamba	Efficient alternative to Transformer with linear complexity. Purpose: Provides faster inference and lower memory usage for ultra-long wireless signal sequences, meeting real-time requirements.
	Multimodal Fusion	Maps vision, text, and sensor data to a unified semantic space (e.g., CLIP [76]/ViT). Purpose: Theoretical basis for the IBSA perception layer, enabling heterogeneous fusion of RF signals and visual images.
Training & Alignment	Instruction Tuning	Fine-tuning via high-quality instruction data. Purpose: Enables IBSA to understand and execute natural language control instructions from network management (e.g., “optimize beams”, “activate jamming”).
	RLHF (Reinforcement Learning from Human Feedback)	Optimizes policy via reward models to align with human values. Purpose: Constrains execution layer decisions to strictly follow safety rules and interference limits while maximizing throughput.
	PEFT / LoRA	Freezes model body, training few parameters. Purpose: Lowers compute barriers, enabling low-cost, rapid adaptation at edge BSs for specific environments without full retraining.
	Federated Learning	Distributed privacy-preserving training uploading only gradients. Purpose: Solves data silos and enables collaborative training while protecting user privacy (e.g., vehicle trajectories).
Inference & Acceleration	MoE	Splits model into expert sub-networks with sparse activation. Purpose: Dynamically invokes modules based on tasks (comm. vs. perception) to reduce latency and energy on resource-constrained BSs.
	FlashAttention	Input/Output-aware exact attention acceleration. Purpose: Boosts GPU inference speed, crucial for ms-level response scenarios like V2X collaborative perception.
	Quantization	Reduces numerical precision (e.g., INT4) to compress models. Purpose: Reduces memory footprint, essential for deploying large models from cloud to edge BSs.
Agents & Decision	In-Context Learning	Few-shot learning without parameter updates. Purpose: Enables rapid adaptation to unseen channel environments or new interference patterns via input examples.
	Chain of Thought	Guides step-by-step reasoning. Purpose: Improves accuracy and interpretability in complex logic tasks (e.g., drone intent analysis, fault diagnosis).
	Embodied Agent	Closed-loop interaction with physical environments. Purpose: Core IBSA concept; empowers BSs to perceive, plan, and execute physical actions (beam adjustment, jamming).
Data & Generation	World Models	Digital twins and simulation of physical environments. Purpose: Generates synthetic data obeying physical laws (e.g., collisions) to address scarcity of extreme scenario data.

tive aggregation schemes (e.g. FedAvgM [84], FedProx [85]) mitigates non-i.i.d. data distribution issues. To further enhance privacy and robustness, techniques such as differential privacy noise, secure multi-party computation, or homomorphic encryption can be introduced [86]. Meanwhile, embodied IBSA must defend against malicious clients and poisoning attacks; robust aggregation methods based on median or trimmed mean, reputation scoring, and anomaly detection mechanisms ensure trustworthy model evolution in open environments [87].

2) Security and Trustworthiness: Adversarial Robustness, Model Evaluation, Accountability

As the “intelligent brain” of critical infrastructure, embodied IBSA must withstand multidimensional attacks: adversarial examples can compromise vision/RF recognition; model inversion may leak sensitive scene data; backdoors can trigger uncontrolled behavior under specific stimuli.

For wireless-domain threats such as CSI tampering, defense strategies include channel-consistency verification across multiple antennas, physical-layer authentication using device-specific RF fingerprints, and anomaly detection

Table 7 Overview of key enabling technologies for the embodied IBSA system.

Category	Key Technology	Role and Significance for Embodied IBSA
Cognitive Core (<i>Perception & Reasoning</i>)	Advanced Perception & Fusion (Sec. 5.1.1)	Fuses external sensor data (camera, lidar) with internal wireless signals to overcome individual sensor limitations (e.g., occlusion, weather). Edge deployment ensures millisecond-level inference for safety-critical tasks.
	Edge LAM Training & Inference (Sec. 5.1.2)	Manages the LAM’s cognitive core on resource-constrained edge devices. Uses PEFT (LoRA) for rapid, low-memory adaptation and MoE to reduce latency via “compute-on-demand”.
	Interpretability (Sec. 5.1.3)	Moves beyond standard methods (SHAP) to explain complex spatiotemporal and multimodal fusion decisions. Provides verifiable causal links (e.g., “why this beam was chosen”) for regulatory compliance and operator trust.
System Orchestration (<i>Architecture</i>)	Cloud-Edge-End Coordination (Sec. 5.2.1)	Implements the LAM across heterogeneous hardware by deploying full models in the cloud, medium models at the edge, and lightweight heads on end devices. Uses model compression (quantization, pruning) and gradient compression for efficient updates.
	Resource Co-optimization (Sec. 5.2.2)	Manages the fundamental trade-off between Perception, Communication, and Computation. Uses network slicing to dynamically allocate resources (GPU, spectrum) to ensure QoS for diverse tasks (e.g., high-priority safety vs. low-priority comms).
Security & Governance (<i>Privacy & Trust</i>)	Federated Learning & Privacy (Sec. 5.3.1)	Enables collaborative IBSA model training across multiple BSs and vehicles without sharing raw, sensitive user data (e.g., trajectories). Enhanced with differential privacy or homomorphic encryption for robust security.
	Security & Trustworthiness (Sec. 5.3.2)	Defends the IBSA against multi-dimensional attacks, including adversarial examples (in vision or RF), model inversion, and backdoors. Ensures accountability through auditable logs for decision traceability.
Foundation (<i>Data Engine</i>)	Digital Twin & Simulation (Sec. 5.4)	Provides a safe, controllable, and reproducible environment for training and validating the IBSA’s perception-cognition-execution loop. Generates synchronized multimodal data (RF, vision, lidar) and bridges the “sim-to-real” gap.

based on statistical deviation from expected channel behavior [88]. For visual adversarial examples, certified defenses such as randomized smoothing, input transformation (e.g., JPEG compression, spatial smoothing), and ensemble-based detection can effectively mitigate perturbation-based attacks [89]. Adversarial training, which augments training data with projected gradient descent perturbations, remains a primary defense to improve model robustness against both modalities.

Regarding privacy protection, differential privacy can be integrated at multiple levels: local differential privacy adds calibrated Laplacian or Gaussian noise to raw sensor readings before upload; central differential privacy applies noise during gradient aggregation in federated learning; and the privacy budget should be carefully tuned to balance utility and protection [86].

For evaluation, comprehensive security benchmarks covering the full perception-cognition-execution chain are

needed, encompassing robustness, transferability, and false positive/negative trade-offs. On accountability, auditable logs, decision-traceability graphs, and compliance interfaces delineate responsibilities between the model and operators to satisfy regulatory and legal requirements.

4. Digital Twin and Simulation Platforms

Training and validating the IBSA’s LAM requires rapid iteration under safe, controllable, and reproducible conditions, a need met by the digital twin and simulation platforms illustrated at the top of Fig. 2. Photo-realistic traffic and vision simulators such as CARLA [36], AirSim [39], Blender [37], together with wireless channel simulators like Sionna [42], Wireless InSite [41], WaveFarer [43], REMCOM, form an end-to-end digital twin platform. By employing a unified scene-description language and data bus, multimodal data (vision, lidar point clouds, radio-frequency signals, positioning) are collected synchronously; domain randomization

and procedural generation help bridge the sim-to-real gap. Achieving this cross-platform temporal alignment is a critical challenge, often addressed using hardware triggers (e.g., precision time protocol) or middleware like the robot operating system to apply unified timestamps to heterogeneous data streams, as demonstrated in practical dataset generation efforts [25]. Furthermore, closed-loop simulation validates the perception-decision-action chain offline, and synthetic data feed back into large-model pre-training and sparse scenario augmentation.

These enabling technologies collectively underpin embodied IBSA's transition from a "large-model system" to a "deployable, governable, large-model-driven agent." They form a closed-loop across computation, data, algorithms, resource management, and security compliance, providing the technical foundation for realizing the integrated perception-communication-intelligence vision of 6G.

VI. Evaluation Metrics and Benchmarks

After examining in the previous section the key technologies that enable embodied IBSA deployment, we now turn to the construction of a holistic evaluation framework designed to quantify the performance of the entire IBSA architecture depicted in Fig. 2, one that spans perception, network, and agent, to ensure the system's behavior in real-world deployments is quantifiable, comparable, and amenable to iteration. This section is organized around four pillars: definitions of metrics, methods of measurement, benchmark resources, and the integrated evaluation protocol. Table 8 summarizes this section, and its main components are detailed as follows.

1. Perception Metrics

The following metrics are from [90].

- **False Positive Rate (FPR)**
FPR = number of false alarms / number of negative samples. Alternatively, for video streams or radar point clouds, report false alarms per unit time or per unit area (e.g., per km²) to characterize the system's tendency to over-alert.
- **False Negative Rate (FNR)**
FNR = number of missed detections / number of positive samples. Note that Recall = 1 - FNR.
- **Precision and Recall**
 $\text{Precision} = \frac{TP}{TP+FP}$, $\text{Recall} = \frac{TP}{TP+FN}$. In detection tasks, report the precision-recall curve and compute the F1 score: $F1 = 2 \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$.
- **Detection Error Tradeoff (DET) Curve**
Plot FPR against FNR across decision thresholds to visually compare algorithmic trade-offs under different operational priorities (e.g., higher safety vs. lower miss cost).
- **Average False-Alarm Distance / Average Miss Duration**

Incorporating target spatial-temporal distribution, define "average false-alarm distance" (mean distance from a false alarm to the nearest true target) or "average miss duration" (time from target appearance to detection) to assess spatiotemporal performance in tracking and monitoring.

- **Mean Average Precision (mAP)**
Compute the area under the precision-recall curve at various intersection over union (IoU) thresholds. Commonly report mAP@0.5 and mAP@0.5:0.95 for multi-class object detection.
- **Spatial Localization Error**
Measure the Euclidean distance between detected object centers and ground-truth in a global coordinate frame. Report root mean square error (RMSE) separately for near-field and far-field to reflect localization accuracy at different ranges.
- **2D/3D IoU**
Evaluate overlap of planar bounding boxes (2D) or volumetric cuboids (3D). For 3D IoU, also account for yaw and pitch errors in voxel fitting.
- **Distance / Velocity / Angle Estimation Error**
For continuously moving targets, report RMSE (and 95% confidence intervals) for distance, speed, and attitude angle estimates, e.g., for multi-step (1 s, 3 s) trajectory prediction tasks.
- **Missing-Rate**
After sensor fusion, the fraction of dropped frames across modalities. Used to evaluate data integrity and robustness under high load or adverse conditions.

2. Network Metrics



The following metrics are from [97].

- **Throughput (Gbps/Hz)**
Data rate per unit bandwidth, reported separately for uplink (sensor data return) and downlink (control command broadcast).
- **End-to-End Latency (ms)**
Round-trip delay from terminal data upload, through base-station inference, to issuance of control commands. Emphasize the 99.9th percentile to capture tail latency.
- **Energy per Inference (J)**
Energy consumed for each fused perception-inference task. Break down total into GPU, RF front end, digital signal processor (DSP), etc.
- **Interference Level**
Quantify the impact of embodied IBSA's RF perception or jamming operations on legacy communications, e.g., via signal-to-interference-plus-noise ratio (SINR) degradation or block error rate (BLER) increase in neighboring cells.

3. Agent-Level Metrics

Table 8 Holistic evaluation framework for embodied IBSA (summary of Section 6).

Evaluation Pillar	Key Metrics / Components	Purpose & Description
Perception Metrics (Sec. 6.1)	FPR, FNR, Precision/Recall (F1), DET Curve, mAP, Localization Error (RMSE), 2D/3D IoU, Velocity/Angle Error, Missing-Rate.	To quantify the IBSA’s accuracy, reliability, and precision in environmental perception and object detection.
Network Metrics (Sec. 6.2)	Throughput (Gbps/Hz), End-to-End Latency (ms, 99.9th percentile), Energy per Inference (J), Interference Level (SINR/BLER degradation).	To measure the communication efficiency, response speed (especially tail latency), resource consumption, and impact on existing network services.
Agent-Level Metrics (Sec. 6.3)	Task Completion Rate (TCR), Robustness Drop (under adverse conditions).	To evaluate the IBSA’s overall effectiveness and resilience (robustness) in executing its complete “perception-cognition-execution” closed-loop tasks.
Benchmark Resources (Sec. 6.4)	Public Datasets: Provide annotated real-world samples (e.g., vision, radar, wireless channels). Digital-Twin Platforms: (e.g., CARLA + Sionna) Generate controllable, synchronized synthetic data and ground-truth.	To provide standardized, reproducible, and controllable testing environments (both real and synthetic) for training, validation, and comparison.
Integrated Protocol (Sec. 6.5)	3-Stage Pipeline: 1. Offline Simulation, 2. Hardware-in-the-Loop (HIL) Testing, 3. Field Testing. Composite Score: IBSA-Score (Weighted aggregation of key metrics).	To define a holistic testing pipeline from lab to real-world, culminating in an application-centric single score for comprehensive comparison.

**Figure 5** Performance comparison of different fusion strategies across modalities (lidar, camera, RF). Data is derived from [25]. The significant gap between “No Fusion” and cooperative methods (e.g., HEAL [91], CoAlign [92], F-Cooper [93], V2X-ViT [94]) demonstrates the necessity of network-level collaboration for the IBSA. AP: average precision; Sync: synchronous; Async: asynchronous; AM: average transmission cost in mega byte).**Figure 6** Ablation study on perception modalities under different difficulty levels (easy, moderate, hard). Data is derived from [25]. The consistent performance of radar (equivalent to wireless data of IBSA) and the superior performance of fusion models highlight the necessity of incorporating wireless sensing for robust perception. M2-Fusion comes from [95]. PP: PointPillars [96].

- **Task Completion Rate (TCR) [98]**

Proportion of closed-loop embodied IBSA processes,

from object detection and beamforming/jamming decisions to coordinated control, that successfully meet

safety and real-time constraints.

- **Robustness Drop**

Performance degradation (in mAP or TCR) under adverse conditions (e.g., fog, occlusion, channel fading) relative to baseline.

4. Benchmark Resources

To ensure repeatability and comparability, the embodied IBSA evaluation framework leverages both public multi-modal datasets and an end-to-end digital-twin simulation platform. Public datasets provide annotated real-world samples (vision, radar, wireless channels) [25], while simulation tools (e.g., CARLA, Sionna) generate controllable, customizable synchronized sensor data and ground-truth trajectories [15]. Together, they comprehensively support perception, network, and agent metrics.

Notably, the V2X-Radar dataset [25] provides critical benchmarks for cooperative perception, offering quantitative evidence for the design principles of the IBSA architecture. As visualized in Fig. 5, comparisons between “No Fusion” and various cooperative strategies demonstrate that single-node perception is insufficient for complex environments, whereas multi-source fusion yields significant performance gains. This validates the necessity of the IBSA’s collaborative network structure. Furthermore, regarding the sensing modality, Fig. 6 illustrates that while optical sensors are susceptible to degradation, the inclusion of wireless sensing data is essential for maintaining robustness under varying difficulty levels. This supports the IBSA’s core premise of utilizing wireless signals for resilient environmental perception.

5. Integrated Evaluation Protocol

To jointly optimize individual agent and overall network metrics, we propose a four-stage testing pipeline with clearly defined terminology and procedures at each step:

- **Offline Simulation**

Generate “standard conditions” (normal traffic, good channels) and “extreme conditions” (harsh weather, high interference) in the digital twin. Perception, communication/network, and computation metrics for pretrained and fine-tuned models to establish baselines [16].

- **Hardware-in-the-Loop Testing (HIL)**

Deploy models to a prototype BS integrating software-defined radio (SDR) and GPU. In a closed-lab setup, emulate real-world links, record inference latency, GPU/RF power curves, and network throughput/interference metrics in real time [15].

- **Field Testing**

Co-deploy embodied IBSA and a traditional evolved node B (eNB) (or other baseline) at the same test site. Compare TCR, 99.9th percentile latency, $\Delta\text{SINR}/\Delta\text{BLER}$ under identical traffic and channel conditions to quantify real-world gains.

- **Composite Embodied IBSA-Score**

To yield a single comprehensive performance indicator, we propose the composite IBSA-Score. We define this score by aggregating N core metrics m_i (detection accuracy, Recall, mAP, 1/center-error, normalized throughput, 1/latency, 1/energy, TCR) with predefined weights: $\text{IBSA-Score} = \sum_{i=1}^N w_i m_i$, $\sum_{i=1}^N w_i = 1$. Weights can be tuned per application scenario (e.g., *safety* > *latency* > *throughput*) to yield a single comprehensive performance indicator. This tuning is intentionally application-centric, reflecting operational priorities, much like setting critical hyperparameters in a machine learning model based on the end goal. While a deep exploration of specific weighting methodologies is outside the scope of this review, formal methods such as the analytic hierarchy process could be employed to systematically derive these weights based on operator-defined pairwise comparisons of the metrics’ importance.

VII. Challenges and Future Directions

Building on the metrics and evaluation protocol introduced above, we now pivot from what to measure to what remains hard, outlining open problems and concrete directions for the embodied IBSA research community.

1. Advancing the Cognitive Core: Robustness, Adaptation, and Action

1) Robust Multimodal Perception Algorithms

Cross-modal alignment under occlusion, glare or spectrum scarcity demands architectures that speak “RF + vision + point-cloud”. Transformers such as MVX-ViT in [17] encode radio tokens alongside image patches and use spatio-temporal attention to fuse beam-space priors with dense semantics, reducing missed detections and improving beam prediction accuracy in V2X scenes. Research frontiers now include causal fusion, calibrated uncertainty and on-device distillation so that perception meets millisecond edge budgets.

2) Few-Shot Adaptation & Lifelong Learning

Real roads and air corridors evolve faster than any static dataset. Parameter-efficient fine-tuning (LoRA, adapters) lets each BS personalise the core LAM with only kilobytes of updates, while federated aggregation guards privacy under non-independent and non-identically distributed traffic. Open problems include catastrophic forgetting, poisoned updates, and balancing exploration-exploitation in continually shifting wireless channels.

3) LAM-Driven Executable Policies via Reinforcement Learning

LAMs can leverage reinforcement learning to translate their rich world models into sequential decision-making capabilities. This allows an agent to convert perception outputs into safe, executable actions, such as beamforming, jamming, or

path-planning, that adhere to strict QoS and legal constraints. Key challenges in operationalizing this approach lie in the effective use of offline RL with telecom digital twins, the design of reward functions that properly balance throughput and safety, and the development of formally verified policy updates for mission-critical links [99].

2. System Realization: Optimization, Energy, and Simulation

1) Joint Perception-Communication-Computation Optimisation

Perception accuracy, spectrum use and GPU cycles are tightly coupled. Cross-layer Lagrangian solvers and graph-RL schedulers have been proposed to minimise 99.9th percentile latency while maximising mAP and throughput, yet real-time feedback between the physical layer and application layers is still missing. Priority pre-emption for bursty events (e.g., rogue drones) and carbon-aware scheduling are fertile grounds for exploration.

2) Energy Modelling and Resource Efficiency

Fine-grained power attribution for RF front-ends and accelerators is largely absent. Accurate joule-per-inference models would unlock dynamic voltage/frequency scaling, smart sleeping of antenna panels and carbon-aware orchestration. Paired with the optimisation in Subsection 7.2.1, this enables longer UAV patrol endurance and greener 6G infrastructure.

3) Digital Twins and Sim-to-Real Transfer

Photo-realistic engines (CARLA [36], AirSim [39]) plus RF ray-tracers (Sionna [42], Wireless InSite [41]) already support closed-loop beam-selection studies, but residual gaps persist in urban-canyon multipath and extreme weather. Domain randomisation, generative replay and hybrid hardware-in-the-loop prototypes are required to guarantee that policies generalise from “silicon city” to real asphalt.

3. Standardization and Evaluation Ecosystem

1) Unified Benchmarks & Open Datasets

A lasting obstacle for embodied integrated perception, communication, and computing and embodied IBSA research is the absence of universally accepted datasets and KPI suites. Public corpora such as DeepSense-6G in [14] already combine mmWave CSI, lidar and camera data, showing how multi-modal ground truth can be collected at scale. A single composite metric, e.g., the embodied IBSA-Score that blends perception, network and agent KPIs, must be elevated from lab script to community leaderboard. Future work should extend benchmarks to low-altitude UAV safety and long-tail driving, release time-synchronised raw RF captures, and publish hardware-in-the-loop test scripts so that academic code and 3rd generation partnership Project (3GPP) conformance reports become mutually reproducible.

A critical goal for these future benchmarks will be to provide quantitative comparisons between the proposed IBSA architecture and existing 5G-advanced ISAC BSs. Although

IBSA may incur higher computational overhead due to its LAM core, its expected performance gains are significant. For instance, higher mAP is expected through robust multi-modal fusion (e.g., RF plus vision) that overcomes the limitations of RF-only sensing in adverse conditions [25]. Furthermore, lower perceived latency is anticipated, enabled by the tight coupling of the perception-cognition-execution loop at the edge, which allows for millisecond-level local actions without core-network round trips. Finally, system-level energy efficiency should be realized via joint perception-communication-computation co-optimization and resource slicing, aiming to minimize the energy-per-task rather than just raw hardware power. Establishing public testbeds to validate these hypotheses, using the IBSA-Score as a holistic metric, remains a crucial next step.

2) Standardisation and Industry Momentum

Scaling from prototypes to city-wide roll-outs demands a single, interoperable playbook. Both 3GPP and ETSI are making critical progress. 3GPP’s study on ISAC (TR 22.837 [100]) is defining shared channel models and perception-first KPIs across bands. Concurrently, ETSI’s industry specification group (ISG) on ISAC and its work on AI-native systems are establishing foundational use cases and architecture principles [101].

From a compatibility perspective, the IBSA architecture aligns with these efforts but introduces specific challenges that require further standardization. For instance, IBSA’s reliance on multimodal fusion (RF, visual, lidar), as discussed in this paper, requires data formats and synchronization protocols that go beyond the RF-centric perception data currently prioritized in TR 22.837. Similarly, the “cognition” layer of IBSA necessitates standardized interface protocols for LAM lifecycle management (like federated updates and MoE routing), which expands on 3GPP’s existing AI for new radio work-item [102].

Therefore, a key future task is to bridge this gap. Aligning academic datasets and the proposed composite embodied IBSA-Score with these emerging 3GPP and ETSI drafts is essential. This will help evolve IBSA from a research concept into a standardized, interoperable component, accelerating multi-vendor interoperability and mass adoption.

4. Trustworthy AI: Explainability, Governance, and Ethics

1) Explainability and Trustworthiness

Explainability is a fundamental requirement, as safety regulators demand a verifiable causal link for every system action [103]. Techniques like self-explainable neural layers and attention heat-maps can clarify how the model aligns radio-frequency and visual data. Furthermore, counter-factual reasoning and signed execution logs are essential for verifying decisions and exploring alternative outcomes. These mechanisms build trust with operators and the public while forming a critical defense against adversarial attacks, such as manipu-

lated RF signals or spoofed visual markers.

2) Ethics, Governance, and Fairness

Massive multi-sensor surveillance raises concerns over privacy, equitable spectrum sharing and low-altitude air-traffic control. Embedding differential-privacy guards, audit trails and human-in-the-loop overrides into network fabric must become a design primitive. Only by coupling technical excellence with transparent governance can 6G embodied agents earn societal acceptance.

VIII. Conclusion

This paper explored how LAMs can enable embodied intelligence in wireless infrastructure by transforming BSs into embodied IBAs. We presented a three-layer perception-cognition-execution architecture that integrates multimodal perception with policy generation and controllable actions, implemented via cloud-edge-end collaboration and parameter-efficient adaptation. Two application cases, cooperative vehicle-road perception for autonomous driving and low-altitude safety against unauthorized UAVs, illustrated how embodied IBAs unlock complementary strengths across radio, vision, and computation to improve both safety and connectivity.

To support rigorous progress, we organized enabling technologies (cognitive core, system architecture, and security as well as trustworthiness) and proposed an evaluation framework that unifies individual agent and overall network metrics. A composite embodied IBA-Score was outlined to facilitate apples-to-apples comparisons and to connect research prototypes with engineering requirements.

Looking ahead, the most pressing needs are: (i) open, time-synchronized multimodal datasets that include raw RF captures and hardware-in-the-loop artifacts; (ii) robust, uncertainty-aware fusion and continual learning that survive occlusion, domain shift, and adversarial conditions; (iii) formally verifiable policy execution that respects QoS, safety, and legal constraints; (iv) joint perception-communication-computation scheduling under tight energy budgets; and (v) standardization of interfaces, metrics, and model lifecycle management, ensuring IBA's multimodal data formats and API protocols align with emerging frameworks from bodies like 3GPP (e.g., TR 22.837) and ETSI ISG-ISAC, coupled with transparent governance. Addressing these challenges will accelerate reproducible, trustworthy deployment of LAM-enabled embodied IBAs and help establish integrated perception, communication, and computation-native intelligence as a practical cornerstone of 6G networks.

References

- [1] X. You, Y. Huang, C. Zhang, *et al.*, "When AI meets sustainable 6G", *Sci. China Inf. Sci.*, vol.68, artilec no.110301, doi:10.1007/s11432-024-4257-6, 2025.
- [2] W. Saad, O. Hashash, C. K. Thomas, C. Chaccour, *et al.*, "Artificial general intelligence (AGI)-native wireless systems: A journey beyond 6G", *Proc. IEEE*, in press, pp.1–39, doi:10.1109/JPROC.2025.3526887, early Access, 2025.
- [3] Y. Huang, "Challenges and opportunities of sub-6 GHz integrated sensing and communications for 5G-Advanced and beyond", *Chin. J. Electron.*, vol.33, no.2, pp.323–325, doi:10.23919/cje.2023.00.251, 2024.
- [4] F. Liu, Y. Cui, C. Masouros, J. Xu, *et al.*, "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, doi:10.1109/JSAC.2022.3156632, 2022.
- [5] X. Fang, C. Lei, W. Feng, Y. Chen, *et al.*, "Sensing-communication-computing-control closed-loop optimization for 6G digital twin-empowered robotic systems", *IEEE J. Sel. Areas Commun.*, vol.43, no.10, pp.3330–3346, doi:10.1109/JSAC.2024.3465133, 2025.
- [6] C. Lei, X. Fang, W. Feng, Y. Chen, *et al.*, "Satellite-UAV networks for 6G control: A sensing-communication-computing-control closed loop perspective", *IEEE Netw.*, vol.39, no.4, pp.62–69, doi:10.1109/MNET.001.2400192, 2025.
- [7] C. Lei, W. Feng, P. Wei, Y. Chen, *et al.*, "Edge information hub: Orchestrating satellites, UAVs, MEC, sensing and communications for 6G closed-loop controls", *IEEE J. Sel. Areas Commun.*, vol.43, no.1, pp.5–20, doi:10.1109/JSAC.2024.3461414, 2025.
- [8] H. L. Cardoso, R. Sousa-Silva, M. Koppinen, and A. Pareja-Lora (Eds.), *Proceedings of the 2nd LUHME Workshop*, LUHME, Bologna, Italy, 2025, ISBN 978-989-9193-73-4, doi:10.21747/978-989-9193-73-4/lan2, <https://aclanthology.org/2025.luhme-1.pdf>.
- [9] C. Zhou, M. Wang, Y. Ma, C. Wu, *et al.*, "From perception to cognition: A survey of vision-language interactive reasoning in multimodal large language models", *arXiv e-Print*, arXiv:2509.25373, 2025.
- [10] A. Liu, W. Jiang, S. Huang, and Z. Feng, "Multi-modal integrated sensing and communication in internet of things with large language models", *IEEE Internet Things J. Mag.*, vol.8, no.5, pp.103–111, doi:10.1109/MIOT.2025.3575888, 2025.
- [11] H. Zou, Q. Zhao, Y. Tian, L. Bariah, *et al.*, "TelecomGPT: A framework to build telecom-specific large language models", *IEEE Trans. Mach. Learn. Commun. Netw.*, vol.3, pp.948–975, doi:10.1109/TMLCN.2025.3593184, 2025.
- [12] F. Zhu, X. Wang, X. Li, M. Zhang, *et al.*, "Wireless large AI model: Shaping the AI-native future of 6G and beyond", *arXiv e-Print*, arXiv:2504.14653, 2025.
- [13] Z. Wang, Y. Shi, Y. Zhou, J. Zhu, and K. B. Letaief, "Edge large AI models: Revolutionizing 6G networks", *IEEE Commun. Mag.*, vol.63, no.10, pp.36–42, doi:10.1109/MCOM.001.2400761, 2025.
- [14] A. Alkhateeb, G. Charan, T. Osman, A. Hredzak, *et al.*, "DeepSense 6G: A large-scale real-world multi-modal sensing and communication dataset", *IEEE Commun. Mag.*, vol.61, no.9, pp.122–128, doi:10.1109/MCOM.006.2200730, 2023.
- [15] W. I. Lab, "Deepverse 6G: A digital twin dataset for wireless communication and sensing", 2025, <https://deepverse6g.net/>.
- [16] K. Huang, S. Mu, J. Jiang, Y. Gao, and S. Xu, "Unreal is all you need: Multimodal ISAC data simulation with only one engine", *arXiv e-Print*, arXiv:2507.08716, 2025.
- [17] G. Gharsallah and G. Kaddoum, "MVX-ViT: Multimodal collaborative perception for 6G V2X network management decisions using vision transformer", *IEEE Open J. Commun. Soc.*, vol.5, pp.5619–5634, doi:10.1109/OJCOMS.2024.3452591, 2024.
- [18] X. Cheng, Z. Huang, L. Bai, H. Zhang, *et al.*, "M3SC: A generic dataset for mixed multi-modal (MMM) sensing and communication integration", *China Commun.*, vol.20, no.11, pp.13–29, doi:10.23919/JCC.fa.2023-0268.202311, 2023.
- [19] X. Cheng, Z. Huang, L. Yu, Yong and Bai, M. Sun, *et al.*, "SynthSoM: A synthetic intelligent multi-modal sensing-communication dataset

- for synesthesia of machines (SoM)", *Sci. Data*, vol.12, article no.819, doi:10.1038/s41597-025-05065-x, 2025.
- [20] B. Salehi, G. Reus-Muns, D. Roy, Z. Wang, *et al.*, "Deep learning on multimodal sensor data at the wireless edge for vehicular network", *IEEE Trans. Veh. Technol.*, vol.71, no.7, pp.7639–7655, doi:10.1109/TVT.2022.3170733, 2022.
 - [21] S. Pradhan, D. Roy, B. Salehi, and K. Chowdhury, "COPILOT: Cooperative perception using lidar for handoffs between road side units", in *Proc. IEEE Conf. Comput. Commun. (INFOCOM)*, pp.1301–1310, doi:10.1109/INFOCOM52122.2024.10621174, 2024.
 - [22] D. Muruganandham, S. Pradhan, J. Gu, T. Braun, *et al.*, "SMART: Sim2Real meta-learning-based training for mmwave beam selection in V2X networks", *IEEE Trans. Mob. Comput.*, vol.24, no.10, pp.1–16, doi:10.1109/TMC.2025.3576203, 2025.
 - [23] S. Kim, J. Moon, J. Kim, Y. Ahn, *et al.*, "Role of sensing and computer vision in 6G wireless communications", *IEEE Wirel. Commun.*, vol.31, no.5, pp.264–271, doi:10.1109/MWC.016.2300526, 2024.
 - [24] F. Bao, X. Wang, S. H. Sureshbabu, G. Sreekumar, *et al.*, "Heat-assisted detection and ranging", *Nature*, vol.619, pp.743–747, doi:10.1038/s41586-023-06174-6, 2023.
 - [25] L. Yang, X. Zhang, J. Li, C. Wang, *et al.*, "V2X-Radar: A multimodal dataset with 4D radar for cooperative perception", *Adv. Neural Inf. Process. Syst. (NeurIPS)*, in press, 2025.
 - [26] T. Mao, L. Liang, J. Yang, H. Ye, *et al.*, "Multimodal-wireless: A large-scale dataset for sensing and communication", in press, no.2511.03220, 2025.
 - [27] G. Charan and A. Alkhateeb, "Sensing-aided 6G drone communications: Real-world datasets and demonstration", *arXiv e-Print*, arXiv:2412.04734, 2024.
 - [28] Y. M. Park, Y. K. Tun, W. Saad, and C. S. Hong, "Resource-efficient beam prediction in mmWave communications with multimodal realistic simulation framework", *arXiv e-Print*, arXiv:2504.05187, 2025.
 - [29] L. Cazzella, F. Linsalata, M. Magarini, M. Matteucci, and U. Spagnolini, "A multi-modal simulation framework to enable digital twin-based V2X communications in dynamic environments", in *2024 IEEE 100th Vehicular Technology Conference (VTC2024-Fall)*, pp.1–6, doi:10.1109/VTC2024-Fall63153.2024.10757947, 2024.
 - [30] G. Gharsallah and G. Kaddoum, "Multimodal collaborative perception for dynamic channel prediction in 6G V2X networks", *IEEE Trans. Mach. Learn. Commun. Netw.*, vol.3, pp.725–743, doi:10.1109/TMLCN.2025.3578577, 2025.
 - [31] B. Salehi, U. Demir, D. Roy, S. Pradhan, *et al.*, "Multiverse at the edge: Interacting real world and digital twins for wireless beamforming", *IEEE/ACM Trans. on Netw.*, vol.32, no.4, pp.3092–3110, doi:10.1109/TNET.2024.3377114, 2024.
 - [32] H. Liu, D. Guo, and A. Cangelosi, "Embodied intelligence: A synergy of morphology, action, perception and learning", *ACM Comput. Surv.*, vol.57, no.7, ISSN 0360-0300, 2025.
 - [33] X. Wang, F. Zhu, Z. Yang, C. Huang, *et al.*, "Bridging physical and digital worlds: Embodied large AI for future wireless systems", *arXiv e-Print*, arXiv:2506.24009, 2025.
 - [34] S. Ruan, R. Wang, X. Shen, H. Liu, *et al.*, "A survey of multi-sensor fusion perception for embodied AI: Background, methods, challenges and prospects", *arXiv e-Print*, arXiv:2506.19769, 2025.
 - [35] Epic Games, "Unreal Engine (version 5.6)", [Software] <https://www.unrealengine.com/>, 2025.
 - [36] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator", in *Proc. Conf. Robot Learn. (CoRL)*, pp.1–16, 2017.
 - [37] Blender Development Team, "Blender (version 4.5)", [Computer software] <https://www.blender.org/>, 2022.
 - [38] Blesor Project, "BlenSor: Blender sensor simulation toolbox", <https://www.blesor.org/>.
 - [39] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles", in *Field and Service Robotics*, 1705.05065, <https://arxiv.org/abs/1705.05065>, 2017.
 - [40] Khronos Group, "OpenGL 4.6 specification", <https://www.khronos.org/registry/OpenGL/specs/gl/glspec46.core.pdf>.
 - [41] Remcom, "Wireless InSite® 3d wireless prediction software (version 4.0)", [Software] <https://www.remcom.com/wireless-insite-propagation-software>, 2025.
 - [42] J. Hoydis, S. Cammerer, F. Ait Aoudia, M. Nimier-David, *et al.*, "Sionna", [Software] <https://nvlabs.github.io/sionna/>, 2022.
 - [43] Remcom, "WaveFarer radar simulation software (version 2.1.0.6)", [Software] <https://www.remcom.com/wavefarer-automotive-radar-software>, 2025.
 - [44] K. Armanious, M. Quach, M. Ulrich, T. Winterling, *et al.*, "Bosch street dataset: A multi-modal dataset with imaging radar for automated driving", *arXiv e-Print*, arXiv:2407.12803, 2024.
 - [45] S. Yao, R. Guan, X. Huang, Z. Li, *et al.*, "Radar-camera fusion for object detection and semantic segmentation in autonomous driving: A comprehensive review", *IEEE Trans. Intell. Veh.*, vol.9, no.1, pp.2094–2128, doi:10.1109/TIV.2023.3307157, 2024.
 - [46] X. Wang, K. Li, and A. Chehri, "Multi-sensor fusion technology for 3D object detection in autonomous driving: A review", *IEEE Trans. Intell. Transp. Syst.*, vol.25, no.2, pp.1148–1165, doi:10.1109/TITS.2023.3317372, 2024.
 - [47] Z. Liu, H. Tang, A. Amini, X. Yang, *et al.*, "BEVFusion: Multi-task multi-sensor fusion with unified bird's-eye view representation", in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, pp.2774–2781, doi:10.1109/ICRA48891.2023.10160968, 2023.
 - [48] H. Liu, J. Liu, G. Jiang, and X. Jin, "MSSF: A 4D radar and camera fusion framework with multi-stage sampling for 3D object detection in autonomous driving", *IEEE Trans. Intell. Transp. Syst.*, vol.26, no.6, pp.8641–8656, doi:10.1109/TITS.2025.3554313, 2025.
 - [49] N. Baumann, M. Baumgartner, E. Ghignone, J. Kühne, *et al.*, "CR3DT: Camera-radar fusion for 3D detection and tracking", in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, pp.4926–4933, doi:10.1109/IROS58592.2024.10801848, 2024.
 - [50] Y. Mei, Z. Gao, Y. Wu, W. Chen, *et al.*, "Compressive sensing-based joint activity and data detection for grant-free massive IoT access", *IEEE Trans. Wireless Commun.*, vol.21, no.3, pp.1851–1869, doi:10.1109/TWC.2021.3107576, 2022.
 - [51] X. Zhou, K. Ying, Z. Gao, Y. Wu, *et al.*, "Active terminal identification, channel estimation, and signal detection for grant-free NOMA-OTFS in LEO satellite internet-of-things", *IEEE Trans. Wireless Commun.*, vol.22, no.4, pp.2847–2866, doi:10.1109/TWC.2022.3214862, 2023.
 - [52] K. Ying, Z. Gao, S. Chen, M. Zhou, *et al.*, "Quasi-synchronous random access for massive MIMO-based LEO satellite constellations", *IEEE J. Sel. Areas Commun.*, vol.41, no.6, pp.1702–1722, doi:10.1109/JSAC.2023.3273699, 2023.
 - [53] Z. Li, Z. Gao, B. Ning, and Z. Wang, "Radiation pattern reconfigurable FAS-empowered interference-resilient uav communication", *IEEE J. Sel. Areas Commun.*, in press, pp.1–1, doi:10.1109/JSAC.2025.3617928, 2025.
 - [54] Z. Li, Z. Gao, and T. Li, "Sensing user's channel and location with terahertz extra-large reconfigurable intelligent surface under hybrid-field beam squint effect", *IEEE J. Sel. Top. Signal Process.*, vol.17, no.4, pp.893–911, doi:10.1109/JSTSP.2023.3278942, 2023.
 - [55] Z. Li, Z. Gao, K. Wang, Y. Mei, *et al.*, "Unauthorized UAV countermeasure for low-altitude economy: Joint communications and jam-

- ming based on MIMO cellular systems”, *IEEE Internet Things J.*, vol.12, no.6, pp.6659–6672, doi:10.1109/JIOT.2024.3491796, 2025.
- [56] Z. Li, S. Tan, Z. Gao, Y. Tao, *et al.*, “Chirp delay-doppler domain modulation: A new paradigm of integrated sensing and communication for autonomous vehicles”, *IEEE Netw.*, vol.39, no.6, pp.119–127, doi:10.1109/MNET.2025.3573832, 2025.
- [57] Z. Gao, X. Zhou, B. Ning, Y. Su, *et al.*, “Integrated location sensing and communication for ultra-massive MIMO with hybrid-field beam-squint effect”, *IEEE J. Sel. Areas Commun.*, vol.43, no.4, pp.1387–1404, doi:10.1109/JSAC.2025.3531551, 2025.
- [58] F. Jiang, C. Pan, L. Dong, K. Wang, *et al.*, “A comprehensive survey of large AI models for future communications: Foundations, applications and challenges”, *arXiv e-Print*, arXiv:2505.03556, 2025.
- [59] M. Abel, I. Ahmad, C. A. Casado, R. Berner, *et al.*, “Large Language Models in the 6G-Enabled Computing Continuum: A White Paper”, *Doctor’s thesis*, University of Oulu, 2025.
- [60] F. Jiang, C. Pan, L. Dong, K. Wang, *et al.*, “From large AI models to agentic AI: A tutorial on future intelligent communications”, *arXiv e-Print*, arXiv:2505.22311, 2025.
- [61] Y. Cui, J. Nie, F. Liu, W. Yuan, *et al.*, “Integrated sensing and communication: Towards multifunctional perceptive network”, *arXiv e-Print*, arXiv:2510.14358, 2025.
- [62] Z. Chen, Z. Zhang, C. Liu, and Z. Xing, “Towards wireless native big AI model: The mission and approach differ from large language model”, *Sci. China Inf. Sci.*, vol.68, no.7, article no.170303, doi:10.1007/s11432-024-4464-8, 2025.
- [63] G. Qin, X. Meng, T. Wen, and B. Cai, “Virtual coupling trains based on multi-agent system under communication delay”, *Chin. J. Electron.*, vol.33, no.6, pp.1545–1554, doi:10.23919/cje.2022.00.253, 2024.
- [64] Z. Gao, D. Mi, C. Jiang, S. Chatzinotas, *et al.*, “Emerging space communication and network technologies for sixth-generation ubiquitous connectivity”, *Space: Science & Technology*, vol.4, article no.0239, doi:10.34133/space.0239, <https://spj.science.org/doi/abs/10.34133/space.0239>, 2024.
- [65] X. Liu, Z. Gao, Z. Wan, Z. Wu, *et al.*, “Toward near-space communication network in the 6G and beyond era”, *Space: Science & Technology*, vol.5, article no.0337, doi:10.34133/space.0337, <https://spj.science.org/doi/abs/10.34133/space.0337>, 2025.
- [66] H. Liu, T. Qin, null, T. Mao, *et al.*, “Near-space communications: The last piece of 6G space-air-ground-sea integrated network puzzle”, *Space: Science & Technology*, vol.4, article no.0176, doi:10.34133/space.0176, <https://spj.science.org/doi/abs/10.34133/space.0176>, 2024.
- [67] Z. Zhang, R. He, M. Yang, X. Zhang, *et al.*, “Non-stationarity characteristics in dynamic vehicular ISAC channels at 28 GHz”, *Chin. J. Electron.*, vol.34, no.1, pp.73–81, doi:10.23919/cje.2024.00.003, 2025.
- [68] MultiX Consortium, “MultiX: Multi-modal 6G experimental infrastructure”, <https://multix-6g.eu/>, 2025.
- [69] Y. Li, J. Zhao, J. Liao, and F. Hu, “Cellular V2X-Based integrated sensing and communication system: Feasibility and performance analysis”, *Chin. J. Electron.*, vol.33, no.4, pp.1104–1116, doi:10.23919/cje.2022.00.340, 2024.
- [70] Y. Tian, J. Wang, Y. Wang, C. Zhao, *et al.*, “Federated vehicular transformers and their federations: Privacy-preserving computing and co-operation for autonomous driving”, *IEEE Trans. Veh. Technol.*, vol.7, no.3, pp.456–465, doi:10.1109/TVT.2022.3197815, 2022.
- [71] N. Yuca, F. Klement, M. Kroll, D. R. H. Spath, and R. S. T. Stuebiger, “A survey on privacy-preserving computing in the automotive domain”, *arXiv e-Print*, arXiv:2508.01798, 2025.
- [72] J. Xia, Y. Liu, and L. Tan, “Joint optimization of trajectory and task offloading for cellular-connected multi-UAV mobile edge computing”, *Chin. J. Electron.*, vol.33, no.3, pp.823–832, doi:10.23919/cje.2022.00.159, 2024.
- [73] H. Zhang, B. Li, J. Huang, C. Song, *et al.*, “A parallel multi-demonstrations generative adversarial imitation learning approach on UAV target tracking decision”, *Chin. J. Electron.*, vol.34, no.4, pp.1185–1198, doi:10.23919/cje.2024.00.082, 2025.
- [74] *Radio Regulations*, edition of 2020 ed., ITU, Geneva, Switzerland, ISBN 978-92-61-31311-8, 2020.
- [75] A. Kutynska and M. Dei, “Legal regulation of the use of drones: Human rights and privacy challenges”, *J. Int. Leg. Comm.*, vol.8, no.1, pp.39–55, doi:10.32612/uw.27201643.2023.8.1.pp.39-55, 2023.
- [76] Z. Yang, G. Chi, C. Wu, H. Liu, *et al.*, “Generative AI meets wireless sensing: Towards wireless foundation model”, *arXiv e-Print*, arXiv:2509.15258, 2025.
- [77] N. Shazeer, A. Mirhoseini, K. Maziarz, A. Davis, *et al.*, “Outrageously large neural networks: The sparsely-gated mixture-of-experts layer”, *arXiv e-Print*, arXiv:1701.06538, 2017.
- [78] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions”, *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol.30, 2017.
- [79] M. Sundararajan, A. Taly, and Q. Yan, “Axiomatic attribution for deep networks”, in *Proc. Int. Conf. Mach. Learn. (ICML)*, PMLR, pp.3319–3328, 2017.
- [80] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, “Deep gradient compression: Reducing the communication bandwidth for distributed training”, *arXiv e-Print*, arXiv:1712.01887, 2017.
- [81] D. Alistarh, T. Hoefer, M. Johansson, N. Konstantinov, *et al.*, “The convergence of sparsified gradient methods”, *Adv. Neural Inf. Process. Syst. (NeurIPS)*, vol.31, 2018.
- [82] X. Cao, Z. Lyu, G. Zhu, J. Xu, *et al.*, “An overview on over-the-air federated edge learning”, *IEEE Wirel. Commun.*, vol.31, no.3, pp.202–210, doi:10.1109/MWC.005.2300016, 2024.
- [83] Y. Chen, J. Hu, J. Zhao, and G. Min, “QoS-aware computation offloading in LEO satellite edge computing for IoT”, *Chin. J. Electron.*, vol.33, no.4, pp.875–885, doi:10.23919/cje.2022.00.412, 2024.
- [84] T.-M. H. Hsu, H. Qi, and M. Brown, “Measuring the effects of non-identical data distribution for federated visual classification”, *arXiv e-Print*, arXiv:1909.06335, 2019.
- [85] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, *et al.*, “Federated optimization in heterogeneous networks”, *Proc. Mach. Learn. Syst. (MLSys)*, vol.2, pp.429–450, 2020.
- [86] Q. Yang, Y. Liu, Y. Cheng, Y. Kang, *et al.*, “Federated machine learning: Concept and applications”, *ACM Trans. Intell. Syst. Technol.*, vol.10, no.2, pp.1–19, 2019.
- [87] X. Cao, M. Fang, J. Liu, and N. Z. Gong, “Fltrust: Byzantine-robust federated learning via trust bootstrapping”, *arXiv e-Print*, arXiv:2012.13995, 2020.
- [88] R. Meng, B. Xu, X. Xu, M. Sun, *et al.*, “A survey of machine learning-based physical-layer authentication in wireless communications”, *J. Netw. and Comput. Appl.*, vol.235, article no.104085, ISSN 1084-8045, doi:<https://doi.org/10.1016/j.jnca.2024.104085>, 2025.
- [89] N. Carlini, F. Tramèr, K. Dvijotham, L. Rice, *et al.*, “(certified!!) adversarial robustness for free!”, in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2023.
- [90] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*, MIT press Cambridge, vol.1, 2016.
- [91] Y. Lu, Y. Hu, Y. Zhong, D. Wang, *et al.*, “An extensible framework for open heterogeneous collaborative perception”, in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2024.
- [92] Y. Lu, Q. Li, B. Liu, M. Dianati, *et al.*, “Robust collaborative 3D object detection in presence of pose errors”, in *Proc. IEEE Int. Conf.*

Robot. Autom. (ICRA), pp.4812–4818, 2023.

- [93] Q. Chen, X. Ma, S. Tang, J. Guo, *et al.*, “F-Cooper: Feature based cooperative perception for autonomous vehicle edge computing system using 3D point clouds”, in *Proc. ACM/IEEE Symp. Edge Comput. (SEC)*, pp.88–100, 2019.
- [94] R. Xu, H. Xiang, Z. Tu, X. Xia, *et al.*, “V2X-ViT: Vehicle-to-everything cooperative perception with vision transformer”, in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, pp.107–124, 2022.
- [95] T. Yin, X. Zhou, and P. Krahenbuhl, “Center-based 3D object detection and tracking”, in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp.11784–11793, 2021.
- [96] A. H. Lang, S. Vora, H. Caesar, L. Zhou, *et al.*, “PointPillars: Fast encoders for object detection from point clouds”, in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, pp.12697–12705, 2019.
- [97] International Telecommunication Union, “Report M.2410-0 (11/2017) - Minimum requirements related to technical performance for IMT-2020 radio interface(s)”, *Report*, No. M.2410-0, ITU-R, 2017, https://www.itu.int/dms_pub/itu-r/opb/rep/R-REP-M.2410-2017-PDF-E.pdf.
- [98] B. Albert and T. Tullis, *Measuring the user experience: Collecting, analyzing, and presenting UX metrics*, Morgan Kaufmann, 2022.
- [99] S. Yue, Y. Deng, G. Wang, J. Ren, and Y. Zhang, “Federated off-line reinforcement learning with proximal policy evaluation”, *Chin. J. Electron.*, vol.33, no.6, pp.1360–1372, doi:10.23919/cje.2023.00.288, 2024.
- [100] 3GPP, “Study on integrated sensing and communication”, *Technical Report*, No. 22.837, 3rd Generation Partnership Project, 2022, <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=4044>.
- [101] ETSI, “Integrated sensing and communications (ISAC); use cases and deployment scenarios”, *Group Report*, No. ETSI GR ISC 001 V1.1.1, ETSI, 2025.
- [102] J. Montojo, “AI/ML for NR air interface”, <https://www.3gpp.org/technologies/ai-ml-nr>, 2022.
- [103] S. Buchanan, D. Pai, P. Wang, and Y. Ma, *Learning Deep Representations of Data Distributions*, Online, <https://ma-lab-berkeley.github.io/deep-representation-learning-book/>, 2025.