

©PIXABAY/GERALT

Embodied AI: From LLMs to World Models

Tongtong Feng, Xin Wang, *Member, IEEE*, Yu-Gang Jiang, *Fellow, IEEE*, and Wenwu Zhu, *Fellow, IEEE*

Abstract

Embodied Artificial Intelligence (AI) is an intelligent system paradigm for achieving Artificial General Intelligence (AGI), serving as the cornerstone for various applications and driving the evolution from cyberspace to physical systems. Recent breakthroughs in Large Language Models (LLMs) and World Models (WMs) have drawn significant attention for embodied AI. On the one hand, LLMs empower embodied AI via semantic reasoning and task decomposition, bringing high-level natural language instructions and low-level natural language actions into embodied cognition. On the other hand, WMs empower embodied AI by building internal representations and future predictions of the external world, facilitating physical law-compliant embodied interactions. As such, this paper comprehensively explores the literature in embodied AI from basics to advances, covering both LLM driven and WM driven works. In particular, we first present the history, key technologies, key components, and hardware systems of embodied AI, as well as discuss its development via looking from unimodal to multimodal angle. We then scrutinize the two burgeoning fields of embodied AI, i.e., embodied AI with LLMs/multimodal LLMs

(MLLMs) and embodied AI with WMs, meticulously delineating their indispensable roles in end-to-end embodied cognition and physical laws-driven embodied interactions. Building upon the above advances, we further share our insights on the necessity of the joint MLLM-WM driven embodied AI architecture, shedding light on its profound significance in enabling complex tasks within physical worlds. In addition, we examine representative applications of embodied AI, demonstrating its wide applicability in real-world scenarios. Last but not least, we point out future research directions of embodied AI that deserve further investigation.

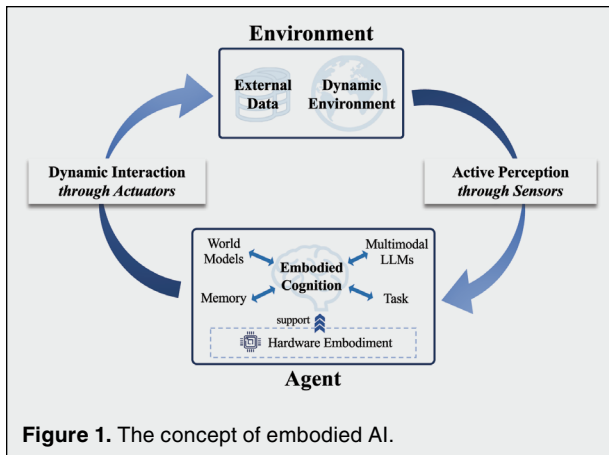
Index Terms—Embodied AI, LLMs, world models.

I. Introduction

Embodied Artificial Intelligence (AI) originated from the Embodied Turing Test by Turing [1], which is designed to explore whether agents can imitate human intelligence to achieve Artificial General Intelligence (AGI). Among them, agents that only solve abstract problems in digital world (cyberspace) are generally defined as disembodied AI, while those that also can interact with the physical world are regarded

Digital Object Identifier 10.1109/MCAS.2025.3603693

Date of current version: 29 December 2025



as embodied AI. Embodied AI builds on foundational insights from cognitive science and neuroscience [2], [3], which claims that intelligence emerges from the dynamic coupling of perception, cognition, and interaction. As shown in Fig. 1, embodied AI includes three key components in a closed-loop manner, i.e., 1) active perception (sensor-driven environmental observation), 2) embodied cognition (historical experience-driven cognition updating), and 3) dynamic interaction (actuator-mediated action control). Besides, hardware embodiment [4], [5], [6] is also critical due to escalating computational and energy demands, particularly under latency and power constraints of devices in real-world deployment scenarios.

The development of embodied AI has evolved from unimodal to multimodal paradigm. In early stage, embodied AI is primarily studied through focusing on individual components with single modality such as vision, language, or action, where the perception, cognition, or interaction component is driven by one sensory input [7], [8], e.g., perception tends to be dominated by the visual modality [9], cognition tends to be dominated by the language modality [10], [11], and interaction tends to be dominated by the action modality [12], [13]. Although these methods perform well within individual components, they are limited by the narrow scope of information provided by each modality and the inherent gaps between modalities across components. The continued development of embodied AI witnesses the limitations of unimodal approaches, promoting a significant shift toward integration of multiple sensory modalities [14],

[15], [16]. As such, multimodal embodied AI [17], [18] naturally arises to create more adaptive, flexible, and robust agents capable of performing complex tasks in dynamic environments.

Large Language Models (LLMs) empower embodied AI via semantic reasoning [19] and task decomposition [20], [21], bringing high-level natural language instructions and low-level natural language actions into embodied cognition. Representative LLM driven works include SayCan [22], which i) provides a real-world pretrained natural language action library to constrain LLMs from proposing infeasible and contextually inappropriate actions; ii) uses LLMs to convert natural language instructions into natural language action sequences; and iii) utilizes value functions to verify the feasibility of natural language action sequences in a particular physical environment. These works suggest that LLMs are extremely useful to robots which aim at acting upon high-level, temporally extended instructions expressed in natural language. However, LLMs are only a part of the entire embodied AI system (e.g., embodied cognition), which is limited by a fixed natural language action library and a specific physical environment, making it difficult for LLM driven embodied AI to achieve adaptive expansion for new robots and environments.

Recent breakthroughs in Multimodal LLMs (MLLMs) [23], [24] and World Models (WMs) [25], [26], [27] have opened up a new frontier in embodied AI research. MLLMs can act on the entire embodied AI system, bridging high-level multimodal inputting and low-level motor action sequences into end-to-end embodied applications. Semantic reasoning [28], [29], [30] leverages MLLMs' cross-modal comprehension to interpret semantics from visual, auditory, or tactile inputs, e.g., identifying objects, inferring spatial relationships, predicting environmental dynamics. Concurrently, task decomposition [31], [32], [33] employs MLLMs' sequential logic to break complex objectives into sub-tasks while dynamically adapting plans based on sensor feedback. However, MLLMs often fail to ground predictions in physics-compliant dynamics [34] and exhibit poor real-time adaptation [35] to environmental feedback.

On the other hand, WMs empower embodied AI by building internal representations [36], [37], [38], [39], [40] and making future predictions [41], [42], [43], [44] of the external world. Such WM driven embodied AI is able to facilitate physical law-compliant embodied

Tongtong Feng is with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China (e-mail: fengtongtong@tsinghua.edu.cn).

Xin Wang and Wenwu Zhu (corresponding author) are with the Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China, and also with the Beijing National Research Center for Information Science and Technology, Beijing 100084, China (e-mail: xin_wang@tsinghua.edu.cn; wwzhu@tsinghua.edu.cn).

Yu-Gang Jiang is with the Institute of Trustworthy Embodied AI, Fudan University, Shanghai 200433, China (e-mail: ygj@fudan.edu.cn).

interactions in dynamic environments. Internal representations compress rich sensory inputs into structured latent spaces, capturing object dynamics, physics laws, and spatial structures, as well as allowing agents to reason about “what exists” and “how things behave” in their surroundings. Simultaneously, future predictions simulate potential rewards of sequence actions across multiple time horizons aligned with physical laws, thereby preempting risky or inefficient behaviors. However, WM driven approaches struggle with open-ended semantic reasoning [45] and lack the ability of generalizable task decomposition [26] without explicit priors.

Building upon the above advances, we further share our insights on the necessity of developing a joint MLLM-WM driven embodied AI architecture, shedding light on its profound significance in enabling complex tasks within physical worlds. MLLMs enable contextual task reasoning but overlook physical constraints, while WMs excel at physics-aware simulation but lack high-level semantics. The joint of MLLM and WM can bridge semantic intelligence with grounded physical interaction. For instance, EvoAgent [46] designs an autonomous-evolving agent with a joint MLLM-WM driven embodied AI architecture, which can autonomously complete various long-horizon tasks across environments through self-planning, self-reflection, and self-control, without human intervention. We believe that designing joint MLLM-WM driven embodied AI architectures will dominate next-generation embodied systems, bridging the gap between specialized AI agents and general physical intelligence.

We summarize the representative applications of embodied AI as service robotics, rescue UAVs, industrial Robots, and others etc., demonstrating its wide applicability in real-world scenarios. We also point out potential future directions of embodied AI, including but not limited to autonomous embodied AI, embodied AI hardware, and swarm embodied AI etc.

As shown in Fig. 2, the rest of this paper is organized as follows. Section II introduces the history, key technologies, key components, and hardware system of embodied AI, discussing the development of embodied AI from unimodal to multimodal angle. Section III presents embodied AI with LLMs/MLLMs, and Section IV presents embodied AI with WMs. Section V introduces our insights on designing a joint MLLM-WM driven embodied AI architecture. Section VI briefly examines applications of embodied AI. Potential future directions are discussed in Section VII.

II. Embodied AI

This section provides a comprehensive overview of embodied AI. We first take a historical view to introduce the development of embodied AI in Section II-A. Based on technological advancements in five foundational areas related to embodied AI, Sections II-B and II-C further review the developmental trajectories of core modules in software algorithms and hardware design, respectively. Finally, Section II-E discusses an overall analysis of the developmental trends from unimodal to multimodal.

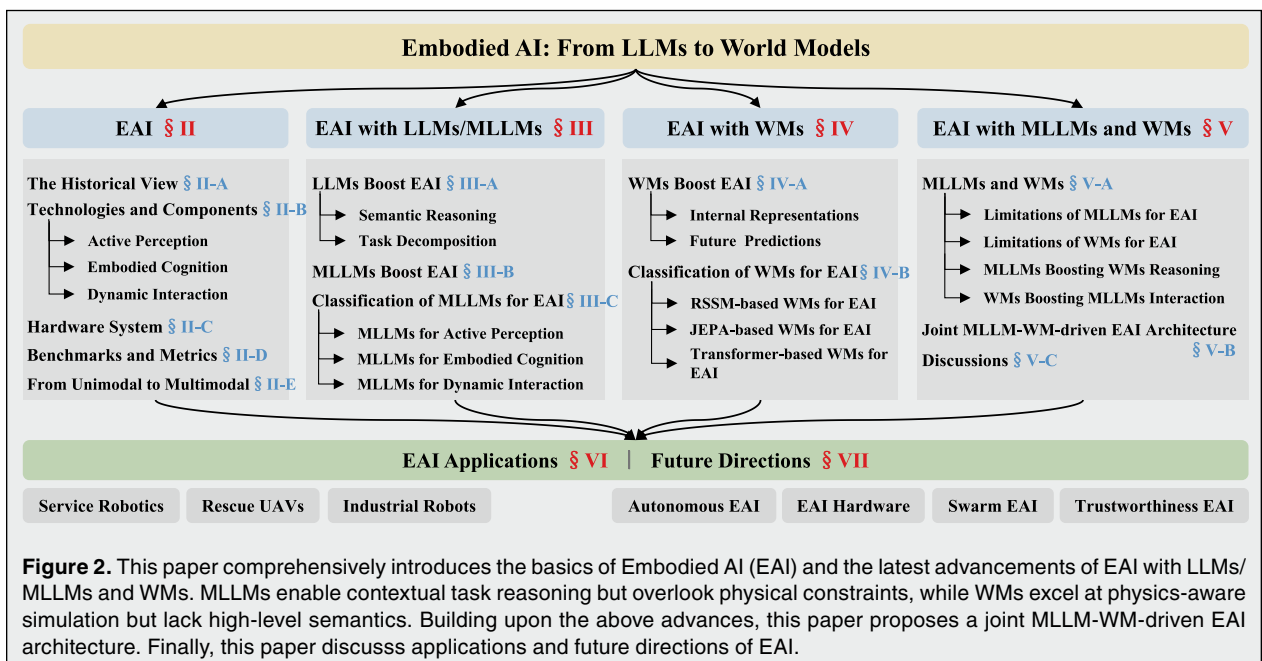


Figure 2. This paper comprehensively introduces the basics of Embodied AI (EAI) and the latest advancements of EAI with LLMs/MLLMs and WMs. MLLMs enable contextual task reasoning but overlook physical constraints, while WMs excel at physics-aware simulation but lack high-level semantics. Building upon the above advances, this paper proposes a joint MLLM-WM-driven EAI architecture. Finally, this paper discusses applications and future directions of EAI.

A. The Historical View

The historical evolution of embodied AI reflects successive transitions from early philosophical foundations to technological breakthroughs in robotics and the rise of learning-driven paradigms, while recent progress in LLMs and WMs is driving an ongoing shift toward the next phase of development.

The theoretical roots of embodied AI trace to 1950, when Turing introduced the foundational idea that intelligence is inherently linked to physical experience [1]. In the 1980s, cognitive science further formalized this view. Lakoff and Johnson emphasized that human cognition arises from bodily experience rather than disembodied symbolic computation [47], while Harnad’s symbol grounding problem highlighted the necessity of connecting symbolic representations to sensory-motor reality [48]. Technological advances in robotics during the late 1980s and 1990s brought these ideas into practice. Brooks proposed the subsumption architecture [49], [50], promoting behavior-based control through layered, reactive modules grounded in sensorimotor loops. The Cog project [51] advanced this line by constructing humanoid robots capable of developmental learning, imitation, and social interaction. Recently, the success of the learning-driven paradigm has driven the shift in embodied AI from motion control of robots to adaptive interaction [52]. In particular, the development of deep learning enables robots to learn complex nonlinear mappings from raw sensor data to action policy, significantly improving navigation and manipulation tasks [53], [54].

While embodied AI has made notable advances, achieving self-reflection intelligence in dynamic, uncertain environments remains a key challenge. Recent progress in LLMs/MLLMs [23], [24] and WMs [25], [26], [27] have progressively shown promise in overcoming these challenges.

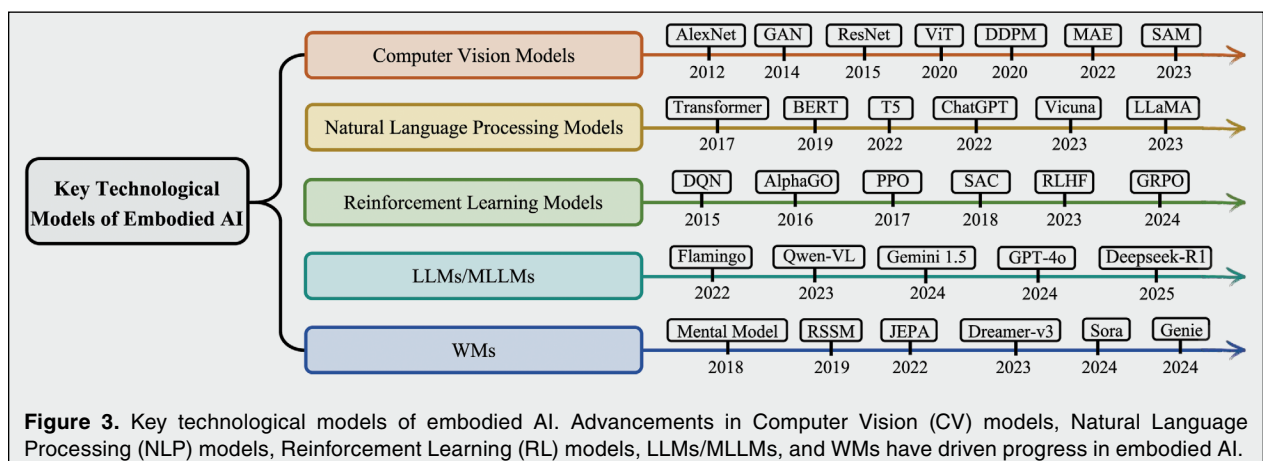
B. The Key Technologies and Components

Before discussing the ongoing changes, we systematically review the development of key technologies and components.

1) *Key Technologies of Embodied AI*: The rapid development of embodied AI is closely tied to advances in foundational technological models such as Computer Vision (CV) models, Natural Language Processing (NLP) models, Reinforcement Learning (RL) models, LLMs/MLLMs, and WMs (as shown in Fig. 3), which can significantly enhance the capabilities of agents in perception, cognition and interaction.

Specifically, Classic models in computer vision, such as AlexNet [55], GAN [56], ResNet [57], ViT [58], DDPM [59], MAE [60], and SAM [61] provide the perceptual foundation for embodied agents to interpret high-dimensional sensory inputs in complex environments. In the field of NLP, the evolution from foundational architectures like Transformer [62], BERT [63], and T5 [64] to large-scale systems such as ChatGPT [65], Vicuna [66], and LLaMA [67], has equipped embodied agents with stronger capabilities in language understanding, task planning, and instruction following. RL offers the core algorithmic framework for agents to learn through interaction with their environments. Representative approaches include DQN [68], AlphaGo [69], PPO [70], SAC [71], RLHF [72], and GRPO [73].

Beyond these classical fields, one of the most promising directions in embodied AI lies in the integration of LLMs/MLLMs with WMs. LLMs and MLLMs (like Flamingo [20], Qwen-VL [74], Gemini-1.5 [75], GPT-4o [76], and Deepseek-R1 [77]) provide agents with the ability to understand instructions, reason over multimodal inputs, and generalize across tasks and environments. In contrast, WMs (like Mental Model [26], RSSM [78], JEPA [27], Dreamerv3 [79], Sora [80], and Genie [36]) enable agents to model and predict environmental



dynamics, supporting imagination-based planning and anticipatory decision-making in dynamic and uncertain environments.

2) *Key Components of Embodied AI*: Driven by advances in these key technologies, embodied AI has experienced rapid progress. In the following, we present a structured overview of developments in three key components.

a) *Active Perception*: Active perception refers to the agent selectively acquiring information from environmental observations [16], [81], [82]. Existing active perception methods can be roughly divided into three categories: visual SLAM, 3D scene understanding, and active environment exploration. To offer an effective perspective on active perception approaches, as summarized in Table 1, we analyze representative methods along three practical dimensions: sensor type, feature type, and applicable scenarios.

Visual SLAM. Simultaneous Localization and Mapping (SLAM) is a pivotal technology enabling agents to both localize themselves and construct environmental maps in unknown environments [9], [83]. As a foundational technology of active perception, visual SLAM has been extensively studied [84], [85]. According to Wang et al. [86], existing methods fall into geometric-based and semantic-based categories. Geometric methods exploit spatial or temporal cues [8], such as dense scene flow [87], [88], triangulation consistency [89], and graph structure [90], [91], performing well in static settings but struggling with dynamic scenes. In contrast, semantic methods improve localization and mapping in dynamic environments by leveraging high-level information. Representative early methods include SLAM++ [92], integrating object-level semantics, and DS-SLAM [93], applying deep learning to dynamic scene understanding. Recent models such as TwistSLAM [94] and GS-SLAM

Table 1. Comparison of three categories of active perception methods including visual SLAM, 3D scene understanding, and active Environment exploration.

Category	Method	Year	Sensor Type	Feature Type	Applicable Scenarios
Visual SLAM	CoSLAM [89]	2012	RGB-D	Geometric + Volumetric	Dynamic SLAM
	SLAM++ [92]	2013	RGB-D	Semantic	Object-level Mapping
	ORB-SLAM [8]	2015	RGB-D + Stereo	Geometric	Dynamic SLAM
	DS-SLAM [93]	2018	RGB-D	Geometric + Semantic	Dynamic SLAM
	TwistSLAM [94]	2022	RGB-D + Stereo	Geometric + Semantic	Dynamic SLAM
	GS-SLAM [95]	2024	RGB-D	Volumetric	Object-level Mapping
	Gaudi [96]	2022	RGB	Volumetric	General Scene Understanding
3D Scene Understanding	Clip2Scene [97]	2023	RGB + Point Cloud	Multimodal	Language-guided Scene Understanding
	OpenScene [98]	2023	RGB + Point Cloud	Multimodal	General Scene Understanding
	Lexicon3D [99]	2024	RGB-D	Semantic	Language-guided Scene Understanding
	GraphDreamer [100]	2024	RGB	Topological + Semantic	Structured Scene Reasoning
	HUGS [101]	2024	RGB-D	Multimodal	General Scene Understanding
	RegionPLC [102]	2024	RGB + Point Cloud	Multimodal	Language-guided Scene Understanding
	MAX [103]	2019	RGB	Semantic	Semantic-guided Exploration
Active Environment Exploration	Active Neural SLAM [104]	2020	RGB-D	Volumetric	Geometry-based Exploration
	APT [105]	2021	RGB	Semantic	Semantic-guided Exploration
	Conan [106]	2023	RGB	Topological	Geometry-based Exploration
	DBMF-BPI [107]	2023	RGB-D	Volumetric	Geometry-based Exploration
	ActiveRIR [108]	2024	RGB + Audio	Multimodal	Cross-modal Active Perception

[95] further enhance robustness by combining geometric optimization with semantic or generative modeling.

3D Scene Understanding. Scene understanding focuses on enabling agents to perceive, segment, and reason about complex environments in a structured and semantically meaningful way. Recent works have advanced this field by integrating vision-language models and generative priors. Early efforts like Gaudi [96] introduced generative models for 3D-aware scene synthesis. Clip2Scene [97] and OpenScene [98] leveraged vision-language embeddings to facilitate label-efficient and open-vocabulary 3D understanding. Structured scene understanding is further enhanced by Lexicon3D [99] and GraphDreamer [100], which model object-level relations in 3D space through structured representations such as scene graphs or semantic lexicons. Meanwhile, region-level multimodal grounding techniques, exemplified by HUGS [101] and RegionPLC [102], incorporate prompts and spatial grounding to achieve fine-grained, goal-conditioned 3D perception. These methods advance holistic, language-aligned 3D understanding.

Active Environment Exploration. Active exploration focuses on enabling agents to autonomously acquire informative observations through interaction with the environment. Early approaches relied on building explicit or implicit environmental models. Representative model-based methods include MAX [103] and Active Neural SLAM [104], which leverage predictive modeling and mapping to support efficient navigation in unseen spaces. In contrast, APT [105] and DBMF-BPI [107] focus on model-free exploration through direct environmental interaction to reduce reliance on explicit modeling. Recent efforts further enhance exploration capabilities by incorporating multimodal perception [108] and semantic reasoning [106].

b) Embodied Cognition: Embodied cognition refers to the emergence of internal representations and reasoning capabilities during the interaction, driven by the agent's self-reflection on its perception and accumulated experience [146, [147], [148]. This component forms the core of embodied AI, enabling agents to perform task planning [149], causal inference [150], and long-horizon reasoning [151], [152]. Recent studies of embodied cognition primarily focus on three aspects: task-driven self-planning, memory-driven self-reflection, and embodied multimodal foundation models. Table 2 presents representative methods analyzed from four perspectives: input modalities, cognition type, reasoning mode, and output type. These dimensions reflect how embodied agents perceive information, form internal models and conduct reasoning.

Task-Driven Self-Planning. In task-driven self-planning, agents autonomously generate structured plans

based on task goals, environmental context, and internal knowledge, without explicit human instructions [153], [154], [155]. Structured learning is a classical solution that develops latent planning spaces or direct policy mappings, achieving high efficiency within training distributions but lacking robustness to out-of-distribution scenarios. Representative approaches include L3P [109], Ego-planer [111], and ETPNav [114]. Recent advances incorporate LLMs or generative models into self-planning. LLM-Planner [110] and AutoAct [112] integrate LLMs into planning by grounding language-guided reasoning into various tasks, while RPG [113] offers a generative perspective, aiming to unify planning and content creation through multimodal reasoning.

Memory-Driven Self-Reflection. Memory-driven self-reflection enables agents to leverage past experiences for long-horizon reasoning, error correction, and self-improvement [46], [156]. Early studies focus on memory processing, including fixed-size replay buffers [157], [158], [159] and differentiable memory architectures [160], [161]. Recent advances introduce reflective mechanisms, where agents summarize or verbalize past experiences to guide future decisions. Reflexion [115] and Reflect [116] enable agents to iteratively self-correct by integrating verbalized feedback into action planning, while RILA [117] extends reflective reasoning to multimodal semantic navigation. Beyond individual reflection, Optimus-1 [118] and REMAC [119] integrate multimodal or multi-agent memory to support long-horizon collaboration. EvoAgent [46] further advances this direction by coupling continual world modeling with a memory-driven planner, enabling fully autonomous evolution across sequential tasks.

Embodied Multimodal Foundation Models. In the era of MLLMs, embodied multimodal foundation models [162], [163], [164] have emerged as one of the most promising solutions for unifying planning, reasoning, and other embodied cognitive capabilities. Recent progress is driven by both data construction and model development. Data efforts focus on constructing high-quality benchmarks to support scalable and cognitively meaningful evaluation, such as MuEP [165], ECBench [166], MFE-ETP [167], and EmbodiedBench [18]. On the model side, recent advances include affordance-grounded agents (e.g., SayCan [120] and GATO [121]) that align language understanding with embodied action spaces, vision-language pretraining approaches (like EmbodiedGPT [122] and Kosmos-2 [123]) that promote scalable embodied reasoning, and object-centric designs (such as MultiPLY [124] and ManipLLM [28]) that enhance manipulation and interaction capabilities. These models collectively aim to build transferable and generalizable embodied AI.

Table 2. Comparison of three categories of embodied cognition methods: task-driven self-planning, memory-driven self-reflection, and embodied multimodal foundation models. I, L and P indicate the image, language and point cloud modalities, respectively.

Category	Method	Year	Input Modalities	Cognition Type	Reasoning Mode	Output
Task-driven Self-planning	L3P [109]	2021	I + L	Planner	Neural + Symbolic	Action
	LLM-Planner [110]	2023	I + L	Planner	Neural + Symbolic	Action
	Egoplaner [111]	2023	I	Planner	Symbolic	Action
	AutoAct [112]	2024	L	Planner	Neural	Action
	RPG [113]	2024	I + L	Planner	Neural	Policy
	ETPNav [114]	2024	I + L	Planner	Neural + Symbolic	Policy
Memory-driven Self-reflection	Reflexion [115]	2023	L	Memory	Beam + Replay	Policy
	Reflect [116]	2023	I + L	Memory	Neural + Symbolic	Policy
	RILA [117]	2024	L	Memory	Neural	Policy
	Optimus-1 [118]	2024	I + L	Memory	Neural	Policy
	EvoAgent [46]	2025	I + L	Memory	Neural	Policy
	REMAC [119]	2025	L	Memory	Neural + Symbolic	Policy
Embodied Multimodal Foundation Models	SayCan [120]	2022	I + L	Planner + Aligner	Neural	Answer + Action
	GATO [121]	2022	I + L + P	Aligner	Neural	Action
	EmbodiedGPT [122]	2023	I + L	Aligner	Neural	Answer + Action
	Kosmos-2 [123]	2023	I + L	Aligner	Neural	Answer
	MultiPLY [124]	2024	I + L	Aligner	Neural	Answer
	ManipLLM [28]	2024	I + L	Aligner	Neural	Answer + Action

c) *Dynamic Interaction*: Dynamic interaction refers to the process in which an agent influences the environment through actions or behaviors grounded in its perception and cognition [168], [169]. Existing research highlights the significance of this capability in enabling agents not only to respond but also to change their surroundings [170], [171]. Studies on dynamic interaction encompass action control, behavioral interaction, and collaborative decision-making. To better understand existing methods, we analyze representative approaches from four perspectives, including input modalities, interaction type, learning paradigm, and task type, as shown in Table 3. These dimensions reflect how agents sense the environment, determine the level and structure of interaction, and generate appropriate behaviors in dynamic multi-agent or human-in-the-loop scenarios.

Action Control. Action control generates motor commands for embodied interaction. Early methods were based on control theory with dynamic system modeling [172], [173] or RL via trial and error [174], [175]. The former is effective for structured or repetitive tasks, while the latter is adaptable to high-dimensional, nonlinear problems. Recent advances mainly follow three directions. Vision-language-action (VLA) models, such as PaLM-E [14],

RT-2 [24], OpenVLA [126], and CogAgent [127], integrate language-guided reasoning for flexible control and have been comprehensively reviewed by Ma et al. [176]. Open-ended frameworks like MineDojo [125] promote continual skill acquisition from open-world knowledge. In addition, Cross-embodiment learning, including CrossFormer [129], HPT [130], and Octo [128], aim to unify policy learning across diverse robots and modalities.

Behavioral Interaction. The behavior of an agent is composed of a sequence of actions. Compared to action control, it emphasizes high-level control through meaningful action patterns, enabling agents to interact in a flexible and goal-directed manner. Recent advances mainly fall into two directions. Imitation learning, including GAIL [131], MGAIL [132], TrafficSim [133], and TrajGen [134], enables efficient acquisition and simulation of complex behaviors. BEHAVIOR-1K [135] provides a large-scale benchmark for evaluating behavior generalization across 1,000 embodied tasks. Behavior-aware enhancement methods, such as AgentLens [136] and ECL [137], improve policy robustness and interpretability. Despite these advances, achieving reliable long-horizon behavioral interaction under sparse feedback remains challenging.

Table 3. Comparison of three categories of dynamic interaction methods including action control, behavioral interaction, and collaborative decision-making, across input modalities, interaction type, modeling paradigm, and task type. I, L, S, P, and T denote Image, Language, State, Proprioception, and Trajectory, respectively. IL denotes Imitation Learning.

Category	Method	Year	Input Modalities	Interaction Type	Learning Paradigm	Task Type
Action Control	MineDojo [125]	2022	I + L	High-level Planning	LLM	Instruction Following
	PaLM-E [14]	2023	I + L + P	Low-level Control	MLLM	Embodied Manipulation
	RT-2 [24]	2023	I + L	Low-level Control	VLA	Embodied Manipulation
	OpenVLA [126]	2024	I + L	Low-level Control	VLA	Embodied Manipulation
	Cogagent [127]	2024	I + L	Low-level Control	MLLM	Instruction Following
	Octo [128]	2024	I + L + P	Low-level Control	VLA	Embodied Manipulation
	CrossFormer [129]	2024	I + L	Low-level Control	VLA	Embodied Manipulation
	HPT [130]	2024	I + L	Low-level Control	VLA	Embodied Manipulation
Behavioral Interaction	GAIL [131]	2016	T	Behavioral	IL	Trajectory Learning
	MGAIL [132]	2017	T	Behavioral	IL	Trajectory Learning
	TrafficSim [133]	2021	T	Behavioral	RL	Trajectory Learning
	TrajGen [134]	2022	I + T	Behavioral	RL	Trajectory Learning
	Behavior-1K [135]	2023	I	Trajectory	IL	Behavior Understanding
	AgentLens [136]	2024	I + S	Trajectory	IL	Behavior Understanding
	ECL [137]	2024	I + L	High-level Planning	IL	Embodied Manipulation
	QMIX [138]	2018	S	Behavioral	RL	Cooperative Decision
Collaborative Decision	Qtran [139]	2019	S	Behavioral	RL	Cooperative Decision
	QPLEX [140]	2019	S	Behavioral	RL	Cooperative Decision
	MAT [141]	2022	S	Behavioral	RL	Cooperative Decision
	CoELA [142]	2024	I + L	Low-level Control	LLM	Cooperative Manipulation
	AgentVerse [143]	2024	L	High-level Planning	LLM	Agent Society Simulation
	MetaGPT [144]	2024	L	High-level Planning	LLM	Agent Society Simulation
	Combo [145]	2024	L	High-level Planning	LLM	Cooperative Planning

Collaborative Decision. Collaborative decision focuses on coordinating multiple agents to achieve shared goals, which is essential for multi-agent systems and human-robot collaboration [177], [178], [179]. Multi-agent RL is a classical solution, with methods like QTRAN [139], QPLEX [140], and Qatten [138] addressing

cooperation via centralized training with decentralized execution. MAT [141] reframes MARL as a sequence modeling problem to mitigate scalability limitations in multi-agent RL. Recent advances integrate LLMs and WMs to enhance multi-agent collaboration. MetaGPT [144], CoELA [142], and AgentVerse [143] leverage LLMs

for task reasoning and coordination, while COMBO [145] composes modular WMs to support scalable collaborative embodied decision.

C. Hardware

As embodied AI evolves, model complexity and size have grown, increasing computational and energy demands. Embodied systems, often operating in dynamic, real-world environments, face strict latency and power constraints—especially at the edge. Thus, developing hardware-friendly directions that maintain performance while optimizing efficiency is crucial for enabling responsive, energy-aware embodied agents. Hardware optimization in embodied AI typically includes four components: hardware-aware model compression, compiler-level optimization, domain-specific accelerators, and hardware-software co-design.

1) *Hardware-Aware Model Compression*: Quantization and pruning [4] are key techniques for reducing model size and computational cost. In embodied agents, which frequently run on low-power embedded hardware, such techniques are vital for enabling fast and efficient inference. Quantization [180] maps weights and activations to lower bit-widths, while pruning [181] removes redundant parameters. To support real-world embodied tasks, such as robotic control or visual navigation, hardware efficiency metrics like power, performance, and area (PPA) can guide bit-width allocation or pruning ratios [182], enabling task-specific trade-offs between accuracy and deployability on physical platforms.

2) *Compiler-Level Optimization*: Compilers bridge high-level embodied AI models and hardware execution. In real-time embodied systems, compiler toolchains are essential for efficient processing of sensor data and decision-making. TVM [5], built on LLVM [183] and CUDA, generates optimized code across platforms. These compilers transform computational graphs through operator fusion and redundant computation elimination [184], enabling responsive behavior. Mapping strategies like loop reordering and tiling enhance data locality, parallelism, and memory access [185], all of which are critical to maintaining low-latency inference in embodied agents.

3) *Domain-Specific Accelerators*: With growing computational demands, domain-specific accelerators (DSAs) are a promising solution for embodied AI. These systems, from robots to AR/VR agents, benefit from fast, energy-efficient hardware tailored for frequent operations. Google's TPU [6], typically integrated with CPUs and GPUs via PCIe, accelerates key operations like matrix multiplication. FPGA-based accelerators [186] allow reconfigurability for adapting to new tasks or changing workloads; CGRA accelerators [187] improve structured,

dataflow-heavy computations common in perception or control. Meanwhile, ASIC-based accelerators [188] offer high throughput and energy efficiency, ideal for deploying high-performance embodied models in real-world environments.

4) *Hardware-Software Co-Design*: Separating algorithm and hardware design can lower runtime efficiency. Hardware-software co-design addresses this through algorithm-system and algorithm-hardware co-optimization. Algorithm-system co-optimization focuses on how to take full advantage of GPU resources like tensor cores and CUDA cores to better support the algorithm [198]. Algorithm-hardware co-optimization aims to improve deployment efficiency by tuning both the model and the hardware architecture. For example, we can perform multi-objective optimization based on the types of operators in the network and the configuration parameters of the hardware [190]. We can also design different numerical quantization schemes along with matching hardware accelerators to better support embodied AI tasks [191].

D. Benchmarks and Evaluation Metrics

Standardized benchmarks and evaluation metrics are crucial for objectively assessing the performance of embodied AI systems. Widely adopted testbeds include Habitat [192], which provides photorealistic 3D indoor environments for navigation and interaction tasks, and ManiSkill [192], offering physics-based manipulation scenarios with diverse object sets. Simulation platforms like MuJoCo [194] enable precise control evaluation in continuous state-spaces, while EmbodiedBench [18] supports holistic evaluation of vision-driven agents across perception, cognition, and interaction. For UAV applications, AirSim [195] provides high-fidelity aerial environments with dynamic obstacles. These testbeds vary in complexity: Habitat excels in visual realism, ManiSkill in object diversity, MuJoCo in physical accuracy, and EmbodiedBench in multimodal integration. Domain-specific benchmarks like BEHAVIOR-1K [135] further enable granular evaluation of 1,000 everyday activities under realistic constraints.

Key evaluation metrics span three critical dimensions: Task Success Rate measures completion accuracy of goal-oriented objectives (e.g., object manipulation or navigation) [24]; Real-time Responsiveness quantifies decision latency and adaptation speed to environmental changes [35]; and Energy Efficiency evaluates computational cost (FLOPS) and power consumption (Watts) during deployment [4]. Additional metrics include Path Length for navigation efficiency [104], Generalization Score for unseen scenarios [196], and Safety Violations for physical compliance [170]. For multi-agent systems,

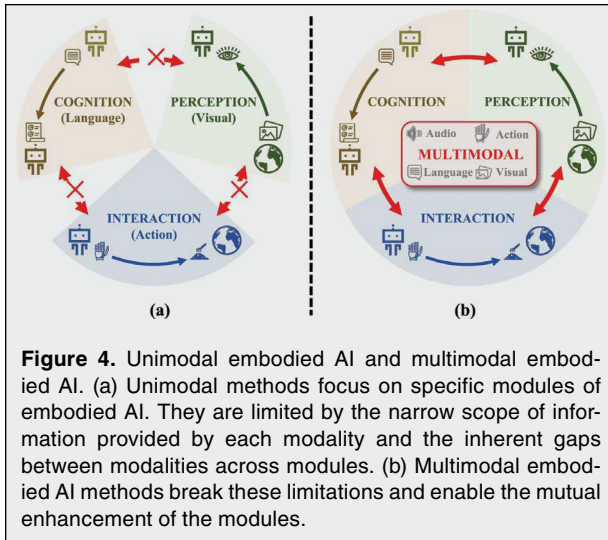


Figure 4. Unimodal embodied AI and multimodal embodied AI. (a) Unimodal methods focus on specific modules of embodied AI. They are limited by the narrow scope of information provided by each modality and the inherent gaps between modalities across modules. (b) Multimodal embodied AI methods break these limitations and enable the mutual enhancement of the modules.

Coordination Efficiency [177] and Communication Overhead [197] provide critical insights. Standardized evaluation protocols like those in MFE-ETP [167] ensure fair cross-modal comparisons, though challenges remain in sim-to-real transfer validation [176].

E. From Unimodal to Multimodal

The development of embodied AI has evolved from unimodal to multimodal systems, as shown in Fig. 4. Initially, embodied AI was primarily concerned with individual modalities, such as vision, language, or action, where the perception, cognition, and interaction were driven by one sensory input [7], [8]. As the field matured, the limitations of unimodal embodied AI became apparent, and there has been a significant shift toward integrating multiple sensory modalities [14], [15], [16]. Multimodal embodied AI is now seen as crucial for creating more adaptive, flexible, and robust agents capable of performing complex tasks in dynamic environments [17], [18].

Unimodal embodied AI has benefited from rapid developments in fundamental areas such as computer vision, natural language processing, and reinforcement learning [12], [13]. These unimodal methods excel in dealing with a specific module in embodied AI. For example, computer vision techniques have driven advances in visual SLAM and 3D scene understanding in the active perception module [9]. Natural language processing techniques, especially LLM, have become popular solutions to address task planning and long-horizon reasoning in the embodied cognition module [10], [11]. Although unimodal embodied AI performs well in independent modules, it always faces two inherent limitations. On the one hand, the information contained in a single modality is limited, hindering the performance of

perception, cognition, and interaction. For example, visual-only systems struggle to understand environments in dynamic or ambiguous settings, while auditory-based systems face challenges in real-world noise and signal processing [17], [198]. On the other hand, diverse and heterogeneous modalities hinder information transfer and sharing among modules. The agent’s perception of the environment fails to facilitate the formation of its cognition, while the evolution of cognition fails to facilitate the interaction with the environment.

In contrast, multimodal embodied AI has emerged as a more promising paradigm [18]. By integrating data from multiple sensing modalities, such as visual, auditory, and olfactory feedback, these methods can provide a more holistic and precise understanding of the environment. More importantly, multimodal embodied AI can facilitate deeper integration among perception, cognition, and interaction. Recent advances in MLLMs and WMs enable agents to more effectively handle multiple modalities, promising to improve the capabilities of embodied AI [44], [80], [199]. The integration of these models is considered a key step toward enabling multimodal embodied AI in dynamic, uncertain environments.

III. Embodied AI With LLMs/MLLMs

This section provides a comprehensive overview of embodied AI with LLMs/MLLMs. We first elaborate in detail how LLMs boost embodied AI in Section III-A and how MLLMs boost embodied AI in Section III-B. Then we discuss the classification of MLLMs for embodied AI in Section III-C.

A. LLMs Boost Embodied AI

LLMs empower embodied AI via semantic reasoning and task decomposition, bringing high-level natural language instructions and low-level natural language actions into embodied cognition.

1) *Semantic Reasoning*: Semantic reasoning [19] leverages LLMs to interpret semantics from text instructions by analyzing linguistic patterns, contextual relationships, and implicit knowledge. Through transformer architectures [62], LLMs map input tokens to latent representations, enabling hierarchical abstraction of meaning across syntactic and pragmatic levels. They employ attention mechanisms to weigh relevant semantic cues while suppressing noise, facilitating logical inference and analogical reasoning. By integrating world knowledge from pretraining corpora with task-specific prompts, LLMs dynamically construct conceptual graphs that align textual inputs with intended outcomes. This process supports multi-hop reasoning

through probabilistic token prediction, resolving ambiguities by evaluating contextual coherence and semantic plausibility.

2) *Task Decomposition*: Task decomposition [20], [21] employs LLMs' sequential logic to break complex objectives into sub-tasks by hierarchically analyzing contextual dependencies and goal alignment. Leveraging chain-of-thought prompting, LLMs iteratively parse instructions into actionable steps, prioritizing interdependencies while resolving ambiguities through semantic coherence checks.

Representative works like SayCan [22] first provides a real-world pretrained natural language actions library, which is used to constrain LLMs to propose both feasible and contextually appropriate actions; then uses LLMs to convert natural language instructions into natural language action sequences; finally uses value functions to verify the feasibility of natural language action sequences in a particular physical environment. These works suggest that LLMs are extremely useful to robots aiming to act upon high-level, temporally extended instructions expressed in natural language. However, LLMs are only a part of the entire embodied AI system, which is limited by a fixed natural language actions library and a specific physical environment, and it is difficult to achieve adaptive expansion in new robots and environments.

B. MLLMs Boost Embodied AI

MLLMs can act on the entire embodied AI system and can solve LLMs' problems well by bridging high-level multimodal inputting and low-level motor action sequences into end-to-end embodied applications (as shown in Fig. 5). Compared with LLMs, semantic

reasoning [28], [29], [30] leverages MLLMs' cross-modal comprehension to interpret semantics from visual, auditory, or tactile inputs, e.g., identifying objects, inferring spatial relationships, or predicting environmental dynamics. Concurrently, task decomposition [31], [32], [33] employs MLLMs' sequential logic to break complex objectives into sub-tasks while dynamically adapting plans based on sensor feedback. MLLMs mainly include Vision-Language Models (VLMs) and Vision-Language-Action models (VLAs).

1) *VLMs for Embodied AI*: VLMs for embodied AI integrate visual and language instruction understanding to enable physical or virtual agents to perceive their environments in goal-driven tasks. Representative works like PaLM-E [14] first train visual and language encodings end-to-end, in conjunction with a pre-trained large language model; then incorporate the results of real-world continuous sensor modalities encodings into VLMs and establish the link between words and percepts; finally, achieve multi-task completion through fixed action space mapping. For navigation, ShapeNet [200], which fine-tunes contrastive embeddings for 3D spatial reasoning, greatly reduces path planning errors. These works suggest that VLMs can combine perception and reasoning in embodied AI to solve a large number of tasks with fixed action spaces.

2) *VLAs for Embodied AI*: VLAs integrate multimodal inputs with low-level action control through differentiable pipelines. Representative works like RT-2 [24] first encode the robot's current image, language instructions, and robot actions at a specific timestep and convert them into text tokens; then use LLMs for semantic reasoning and task decomposition; finally, de-tokenizes generated tokens into the final action. Octo [128]

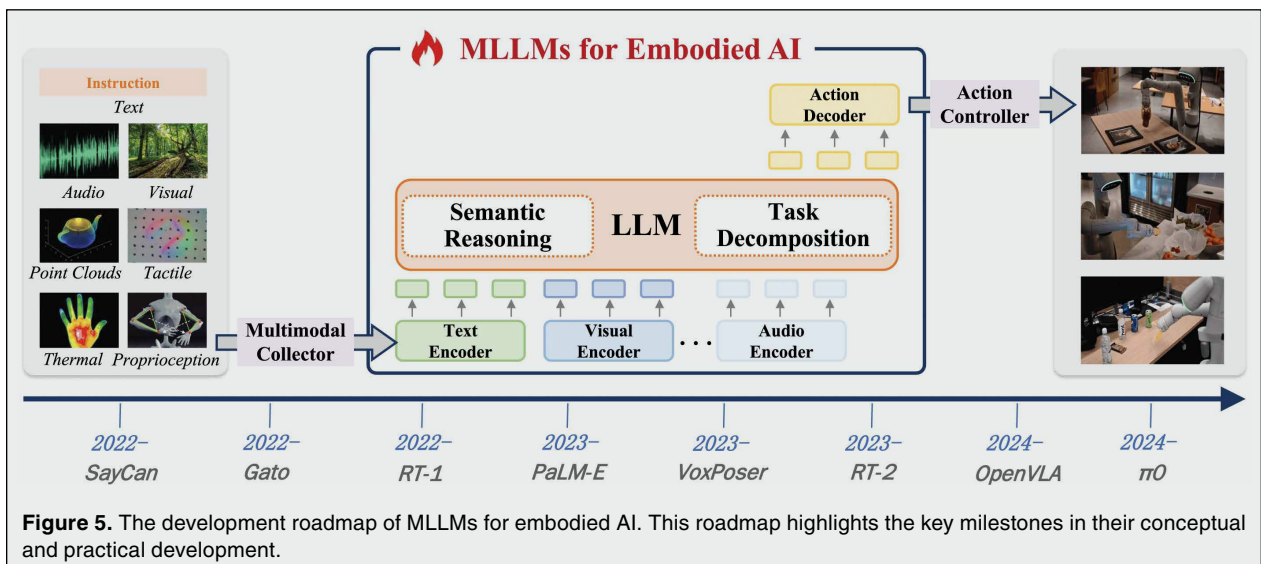


Figure 5. The development roadmap of MLLMs for embodied AI. This roadmap highlights the key milestones in their conceptual and practical development.

pretrains on 100K robot demonstrations with language annotations, achieving cross-embodiment tool use. For dexterous manipulation, PerAct [201] utilizes 3D voxel representations to reach millimeter-level grasp accuracy. These works suggest that VLAs can act on the entire embodied AI system and achieve adaptive expansion in new robots and environments.

C. Classification of MLLMs for Embodied AI

MLLMs can empower active perception, embodied cognition, and dynamic interaction of embodied AI.

1) *MLLMs for Active Perception*: First, MLLMs can enhance 3D SLAM. By grounding visual observations into semantic representations, MLLMs augment traditional SLAM pipelines with high-level contextual information such as object categories, spatial relations, and scene semantics. Representative works like SEO-SLAM [202] utilize MLLMs to generate more specific and descriptive labels for objects, while dynamically updating a multiclass confusion matrix to mitigate biases in object detection. Second, MLLMs can enhance 3D scene understanding. Camera-based perception [30] remains the dominant setup in MLLM-driven embodied AI, as RGB inputs align naturally with the visual-language pretraining of many foundation models [203]. Representative works like EmbodiedGPT [122] leverage this synergy to map 2D visual inputs into semantically rich features aligned with language-based goals. Finally, MLLMs can enhance active environment exploration. MLLMs have also revolutionized how robots interact with their environments, particularly in feedback-driven closed-loop interactions. Representative works like LLM³ [201] focus on structured motion-level feedback, which incorporates signals such as collision detections into the planning loop, allowing the model to iteratively revise symbolic action sequences. MART [204], on the other hand, leverages interaction feedback to improve retrieval quality.

2) *MLLMs for Embodied Cognition*: First, MLLMs can enhance task-driven self-planning. Embodied agents with MLLMs can either directly map high-level goals to structured action sequences [31], or adopt an intermediate planning strategy that continually interacts with the environment to refine their plans [32]. Representative works like CoT-VLA [33] predict intermediate subgoal images that depict the desired outcomes of subtasks, helping the agent visualize and reason through each step of a complex task. Second, MLLMs can enhance memory-driven self-reflecting. MLLMs allow agents to learn from experience using this inherent memory module [128]. Representative works like Reflexion [115] enhance agent performance through self-generated linguistic feedback, which is stored in an episodic memory

buffer and leveraged to guide future planning. Finally, MLLMs can enhance embodied multimodal foundation models. MLLMs can be adapted to the physical world through continued pretraining or fine-tuning in embodied settings. Representative works include QwenVL [74] and InternVL [205], along with models supporting broader modality alignment, such as Qwen2.5-Omni [206].

3) *MLLMs for Dynamic Interaction*: First, MLLMs can enhance action control. MLLMs have ability to decompose complex tasks into actionable subtasks [32]. To further produce continuous control signals for each subtask, MLLMs either generate actions autoregressively in a sequential manner [126], [207] or employ auxiliary policy heads to further process their internal representations [128]. Recent advances also explore generating executable code with MLLMs [208], enabling robots to follow interpretable and adaptable control policies. Second, MLLMs can enhance behavioral interaction. Through interaction with the environment, MLLMs are also capable of generating sequences of behavioral actions in a single step. Representative works like π -0 [31] combine a vision-language backbone with a flow-matching decoder to produce smooth, temporally extended behavioral trajectories. Finally, MLLMs can enhance collaborative decision-making. One line of research focuses on multi-agent systems that aim to achieve human-level coordination and adapt rapidly to unforeseen challenges [209]. For instance, Combo [145] introduces a novel framework that enhances cooperation among decentralized agents operating solely with egocentric visual observations. Other efforts investigate human-agent collaboration. VLAS [210] exemplifies this by aligning human verbal commands with visual context via a speech encoder and a LLaVA-style MLLM [211], enabling fluid and conversational human-agent interaction.

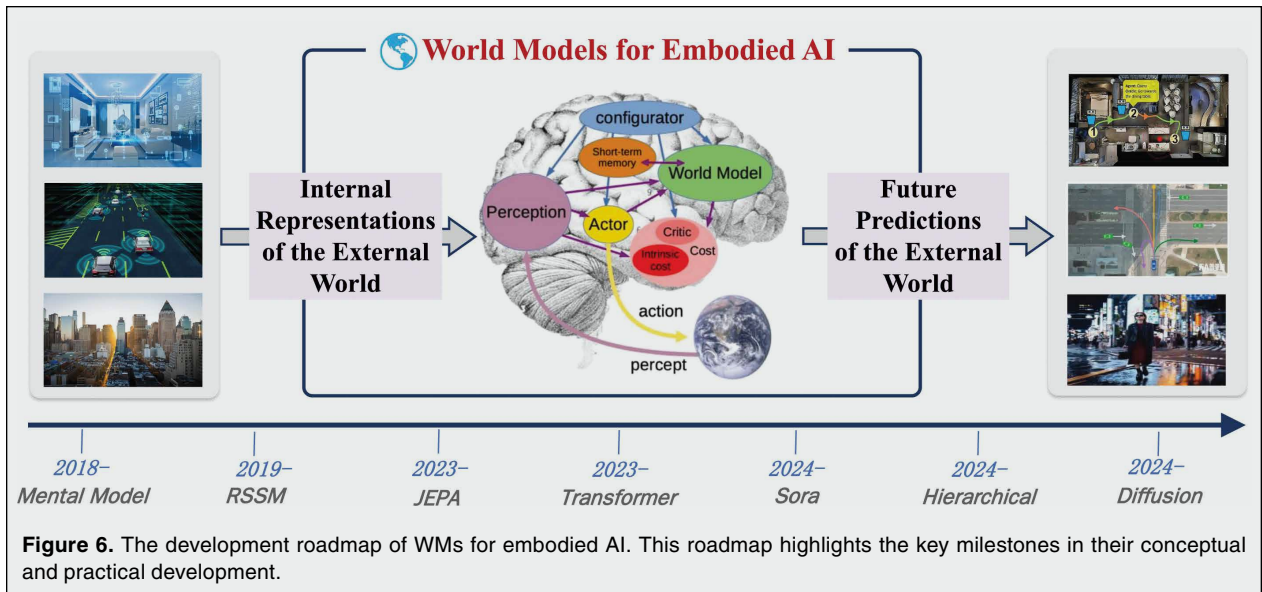
IV. Embodied AI With World Models

This section provides a comprehensive overview of embodied AI with WMs. We first elaborate in detail how WMs boost embodied AI in Section IV-A. Then we discuss the classification of WMs for embodied AI in Section IV-B.

A. World Models Boost Embodied AI

WMs empower embodied AI by building internal representations and future predictions of the external world (as shown in Fig. 6), facilitating physical law-compliant embodied interactions in dynamic environments.

1) *Internal Representations of the External World*: Internal representations compress rich sensory inputs



into structured latent spaces, capturing object dynamics, physics laws, and spatial structures, allowing agents to reason about “what exists” and “how things behave” in their surroundings. These latent embeddings preserve hierarchical relationships [212] between entities and environments, mirroring the compositional nature of reality itself. The structured nature of these representations facilitates generalization across environments, as abstracted principles (like gravity or object permanence) transcend specific instances. Moreover, they support counterfactual reasoning [40] by maintaining disentangled variables for objects’ intrinsic properties [38] and extrinsic relations [39], enabling flexible mental manipulation of individual components. This disentanglement also enhances sample efficiency in learning, as agents transfer knowledge between tasks, sharing latent factors. World models with rich internal representations, can introspect on their own uncertainty about environmental states and actively seek information to resolve ambiguities. By encoding temporal continuity and spatial topology [36], these models naturally enforce consistency constraints during planning, filtering physically implausible actions before execution. Ultimately, such structured latent spaces act as cognitive scaffolding for building causal understanding [37], mirroring how humans develop intuitive theories about their world through compressed sensory experiences.

2) *Future Predictions of the External World:* Future predictions simulate potential rewards of sequence actions across multiple time horizons aligned with physical laws, thereby preempting risky or inefficient behaviors [41], [42]. This predictive capacity bridges short-term actions with long-term goals [43], filtering

out trajectories violating physical plausibility (e.g., walking through walls) or strategic coherence (e.g., depleting resources prematurely). Long-horizon prediction [44] allows adaptive balancing of exploration-exploitation tradeoffs, simulating distant outcomes to avoid local optima while maintaining focus on actionable near-term steps. Crucially, these predictions incorporate uncertainty quantification [41], [213], distinguishing predictable regularities (daily patterns) from stochastic events (sudden changes) to optimize risk-aware planning. The simulation prediction improves sample efficiency [39], [214], [215], [216] by replacing costly trial-and-error with mental rehearsal, particularly valuable in safety-critical domains like autonomous driving or robotic surgery. Furthermore, continuous prediction-error minimization drives iterative model refinement [169], [217], [218], [219], creating self-correcting systems that align their internal physics simulators with observed reality. Such anticipatory capabilities ultimately grant artificial agents human-like foresight, transforming reactive responses into purposeful, future-optimized behaviors.

B. Classification of World Models for Embodied AI

Embodied AI with WMs can mainly be divided into three critical structures: the Recurrent State Space Model-based (RSSM-based) WMs for embodied AI, the Joint-Embedding Predictive Architecture-based (JEPA-based) WMs for embodied AI, and the Transformer-based WMs for embodied AI. Hierarchical-based WMs [220] and diffusion-based WMs [221] are similar to other structures and are shown in Fig. 6.

1) *RSSM-based WMs for Embodied AI:* RSSM constitutes the fundamental architecture underpinning the

Dreamer algorithm family [41], [42], [43], [44]. This framework enhances predictive capabilities in latent representations by acquiring temporal environment dynamics through visual inputs, subsequently enabling action selection via latent trajectory optimization. Through orthogonal decomposition of hidden states into probabilistic and deterministic components, the architecture explicitly accounts for both systematic patterns and environmental uncertainties. Its demonstrated effectiveness in robotic motion control applications has inspired numerous derivative studies building upon its theoretical framework.

2) *JEPA-Based WMs for Embodied AI*: JEPA [27] provides a structure for developing autonomous machine intelligence systems. This architecture establishes mapping relationships between input data and anticipated outcomes through representation learning. Diverging from conventional generative approaches, JEPA operates in abstract latent spaces rather than producing pixel-wise reconstructions, thereby prioritizing semantic feature extraction over low-level signal synthesis. A key methodological foundation of JEPA [213] involves self-supervised training paradigms where neural networks learn to infer occluded or unobserved data segments. Such pretraining on extensive unlabeled datasets enables transfer learning across downstream applications, demonstrating enhanced generalization capabilities for both visual [222], [223] and non-visual domains [224].

3) *Transformer-Based WMs for Embodied AI*: Originating in natural language processing research, the Transformer structure [62] fundamentally relies on attention mechanisms to process input sequences through parallelized context weighting. This design allows simultaneous computation of interelement dependencies, overcoming the sequential processing constraints inherent in Recurrent Neural Networks (RNNs). Empirical evidence demonstrates superior performance in domains requiring persistent memory retention and explicit memory addressing for cognitive reasoning [225], which has propelled its adoption in reinforcement learning research since 2020. Existing advancements have successfully implemented WMs using Transformer variants [38], [40], [226], outperforming RSSM architectures in memory-intensive interactive scenarios [37]. Notably, Google's Genie framework [36] employs the Spatial-Temporal Transformer (ST-Transformer) [227] to create synthetic interactive environments through large-scale self-supervised video pretraining. This breakthrough establishes novel paradigms for actionable world modeling, revealing transformative potential for WMs development trajectories.

V. Embodied AI With MLLMs and WMs

This section provides a comprehensive overview of embodied AI with MLLMs and WMs. We first elaborate in detail on the limitations of MLLMs and WMs for embodied AI and explain how MLLMs boost WMs reasoning, and how WMs boost MLLMs interaction in Section V-A. Then we design a joint MLLM-WM-driven embodied AI architecture in Section V-B. Finally, we discuss the advantages and challenges of new architecture in Section V-C.

A. MLLMs and WMs

MLLMs enable contextual task reasoning but overlook physical constraints, while WMs excel at physics-aware simulation but lack high-level semantics. Their joint bridges semantic intelligence with grounded physical interaction.

1) *The Limitations of MLLMs for Embodied AI (Without WMs)*: MLLMs exhibit two critical limitations in embodied AI applications. First, they often fail to ground predictions [34] in physics-compliant dynamics, leading to impractical plans. For example, ignoring friction or material properties when manipulating objects may cause slippage or task failure. Second, their poor real-time adaptation to environmental feedback limits responsiveness [35]. While MLLMs excel at semantic task decomposition, they struggle to adaptively adjust actions when the environment changes dramatically. These limitations stem from their reliance on static, pre-trained knowledge rather than continuous physical interaction.

2) *The Limitations of WMs for Embodied AI (Without LLMs/MLLMs)*: WMs face limitations in abstract reasoning and generalization. They struggle with open-ended semantic tasks [45] due to their focus on physical simulation rather than contextual understanding. Additionally, WMs lack generalizable task decomposition [26] without explicit priors. For example, a WM model trained on rigid-object manipulation may fail to adapt to deformable materials without extensive retraining. Their predictive accuracy heavily depends on domain-specific interaction records, hindering scalability across diverse environments.

3) *MLLMs Boosting WMs Reasoning*: By leveraging cross-modal alignment and semantic grounding, MLLMs enable WMs to process complex environments dynamically, improving semantic reasoning, task decomposition, and human-robot interaction. 1) MLLMs can enrich WMs by fusing visual, auditory, and textual data into unified semantic representations. For instance, CLIP-based architectures [228] enable agents to align visual scenes with linguistic cues, reducing ambiguity in object recognition [229]. 2) MLLMs can

WMs can play a pivotal role in refining MLLMs by providing physical laws, spatio-temporal relationships, and closed-loop interaction experiences.

augment WM's task decomposition capacity by decomposing high-level goals into executable sub-tasks. Models like GPT-4V [65] generate step-by-step plans using environmental context stored in WM. For robotic manipulation, Code-as-Policies [230] translates natural language instructions into code snippets, leveraging WM to track intermediate states. 3) MLLMs enable WMs to refine internal representations through human feedback. Techniques like Reinforcement Learning with Human Feedback (RLHF) [72] allow agents to update WM priors based on corrective inputs [115]. Those works in this section are all possible ways for MLLMs to boost WMs reasoning, which is not achieved in existing works. 4) *WMs Boosting MLLMs Interaction*: WMs can play a pivotal role in refining MLLMs by providing physical laws, spatio-temporal relationships, and closed-loop interaction experiences. WMs can mitigate MLLMs' inherent limitations in temporal coherence and environmental grounding, enabling more robust decision-making in dynamic embodied tasks. 1) WMs can provide MLLMs with explicit representations of physical laws (e.g., gravity, friction) and commonsense rules to constrain action proposals. For example, Physion++ [231] integrating WM-stored biomechanical models can be used to filter MLLM-generated robotic motions violating torque limits; RoboGuide [232] injects spatial occupancy maps into MLLM planners, preventing collisions during navigation. 2) WMs can stabilize MLLMs reasoning by maintaining spatio-temporal context during multimodal

processing. For instance, MemPrompt [233] can use WM buffers to align visual object trajectories with linguistic descriptions, resolving ambiguities in cluttered environments; RoboMem [234] can leverage WM-prioritized attention to filter irrelevant sensory noise, improving MLLM-based scene understanding. 3) WMs can enable iterative refinement of MLLM outputs through closed-loop interaction. Reflexion [115] can store task-execution histories in WM, allowing MLLMs to correct kinematic errors using failure patterns [230]. Those works in this section are all possible ways for WMs to boost MLLMs' decisions, which has not been achieved in existing works.

B. Joint MLLM-WM-Driven Embodied AI Architecture

We propose a joint MLLM-WM-driven embodied AI architecture (as shown in Fig. 7), shedding light on their profound significance in enabling complex tasks within physical worlds. The specific workflow is as follows, with arrows highlighting the data exchange process.

1) *Robots → Self-State Inputting → MLLMs/WMs → Hardware Embodiment → Robots*: The process initiates with self-state inputting tracking proprioceptive metrics, such as degrees of freedom, number of sensors, etc. These metrics feed into both WMs and MLLMs: WMs use them to build internal representations of the agent's physical state, while MLLMs contextualize these states for task alignment. Hardware embodiment is focused on implementing WMs and MLLMs into physical

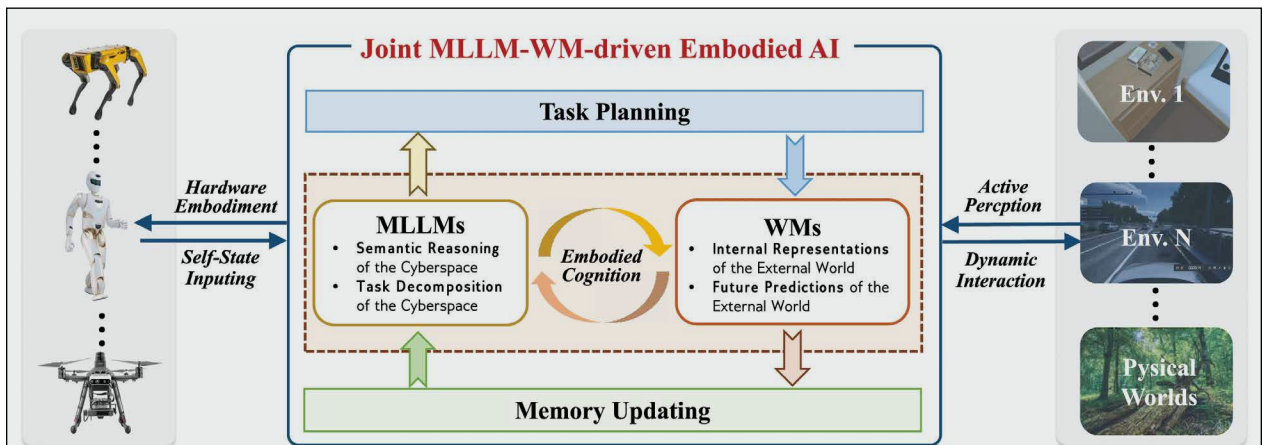


Figure 7. Embodied AI with MLLMs and WMs. MLLMs can enhance WMs by injecting semantic knowledge for task decomposition and long-horizon reasoning, while WMs can assist MLLMs by building the physical world's internal representations and future predictions, making joint MLLM-WM a promising architecture for embodied systems.

MLLMs excel in semantic reasoning, enabling high-level task decomposition, contextual understanding, and adaptive planning by leveraging multimodal inputs.

devices to solve sim-to-real problems. This bidirectional flow ensures actions respect both mechanical limits and high-level goals.

2) *MLLMs* → *Task Planning* → *WMs* → *Memory Updating* → *MLLMs*: MLLMs decompose abstract instructions into subtasks. A forward arrow delivers this plan to WMs, which predict outcomes based on existing environmental modeling. During execution, WMs log outcomes into memory. A vertical arrow transmits these logs to memory updating modules, which structure memory into experiences, represent the forgetting of past task memories, the renewal of current task memories, and the prediction of future task memories. These are then fed back to MLLMs via an arrow, enriching their knowledge base. This enables lifelong learning, where past failures directly inform future planning.

3) *Environments* → *Active Perception* → *MLLMs/WMs* → *Dynamic Interaction* → *Environments*: WMs first drive active perception by predicting key environmental changes. Multimodal inputs are then used to construct

an internal representation of the external world through WMs and semantic reasoning through MLLMs. Then, the task decomposition of MLLMs and future prediction of WMs enable action selection and environmental interaction. Adaptive perception and interaction of dynamic environments are achieved through continuous iteration.

C. Discussions

Joint MLLM-WM offer a promising architecture for embodied AI. As shown in Table 4, MLLMs excel in semantic reasoning, enabling high-level task decomposition, contextual understanding, and adaptive planning by leveraging multimodal inputs. Meanwhile, WMs provide grounded, physics-based simulations of environments, ensuring actions align with real-world constraints. This synergy allows agents to balance abstract reasoning with real-time physical interactions, enhancing decision-making in dynamic settings. For instance, MLLMs can generate task plans while WMs validate feasibility, enabling iterative refinement. Additionally, joint architectures support

Table 4. Qualitative comparison of LLM-only, WM-only, and joint LLM-WM architectures in embodied AI. Low, Medium, High.

Performance	LLM/MLLM-only	WM-only	Joint LLM-WM
Semantic Understanding	Advantages in contextual task reasoning and natural language understanding	Limited in open-ended semantic understanding	Combines high-level semantic abstraction with grounded contextual alignment
Task Decomposition	Sequential logic enables sub-task planning via language prompts	Lacks generalizable task decomposition mechanisms	Semantic plans refined through physical feasibility via joint planning-execution loop
Physics Compliance	Ignores physical constraints and dynamics in real-world interaction	Physics-aware simulation with temporal consistency	Enforces semantic-physical alignment for safe and executable plans
Future Prediction	Lacks imagination-based reasoning	Long-horizon multi-step prediction with uncertainty modeling	Combines symbolic foresight and physically grounded imagination
Real-time Interaction	Poor responsiveness to environmental feedback and significant reasoning latency	Supports real-time predictive control via future state simulation	Enables online adaptation through iterative plan refinement and memory updating
Memory Structure	Sparse and unstructured memory	Structured latent space encodes object dynamics and causal relationships	Integrates semantic memory and world modeling for lifelong learning and reflection
Scalability	Limited to pre-trained task space	Poor transfer to unseen tasks without retraining	Cross-task, cross-domain generalization through symbolic and sensorimotor synergy

The WM maintains an evolving environment model for planning and simulation, while the MLLM grounds commands like “clean up the living room” into adaptive subtasks.

cross-modal generalization, improving robustness in partially observable or novel scenarios by bridging symbolic knowledge and sensorimotor experiences.

The challenges of joint MLLM-WM-driven embodied AI architecture include 1) real-time synchronization between MLLMs’ high-latency semantic processing and WMs’ physics-based representation, often leading to delayed responses in dynamic environments; 2) semantic-physical misalignment, where MLLM-generated plans violate unmodeled physical constraints; and 3) scalable memory management, as continuous updates to WM’s internal states risk overwhelming MLLMs with irrelevant context. Additionally, training such systems requires vast multimodal datasets covering rare edge cases, while ensuring robustness against sensor noise and partial observability remains unsolved. These challenges need lightweight MLLMs inference, tighter feedback loops, and dynamic context-filtering mechanisms to minimize latency.

VI. Embodied AI Applications

This section overviews the application of embodied AI in service robots, rescue robots, and other domains, highlighting trends in joint MLLMs and WMs to advance active perception, embodied cognition and dynamic interaction.

A. Service Robotics

Embodied AI is becoming an important technology in the service field. It helps service robots go beyond fixed rules and perform tasks in a flexible way using different types of information. Recent research highlights its flexible applications across various fields. In domestic settings, systems such as RT-2 [207] and SayCan [120] combine language instructions with robot control, allowing robots to do tasks such as stacking dishes or cooking. Few-shot learning methods like AED [235] acquire new skills from limited demonstrations. In healthcare, robots with multiple types of input can help with reminders, rehabilitation, and companionship. [236], [237]. In public environments, platforms like Habitat [192] and RT-X [238] support navigation and item delivery, even in changing environments, without needing special training for each task. This makes the system more general and useful in real life.

However, current approaches remain limited in handling long-horizon tasks. As illustrated in Fig. 7, the joint

of WMs and MLLMs is emerging as a key strategy for enhancing the autonomy and long-term reasoning capabilities of service robots. The WM maintains an evolving environment model for planning and simulation, while the MLLM grounds commands like “clean up the living room” into adaptive subtasks. This collaboration supports flexible reasoning, goal adaptation, and robust real-world execution.

B. Rescue UAVs

Embodied AI technology is changing the way drones are used in disaster situations. Traditional drones are either manually controlled or rely on pre-built maps when in use, which leads to their inability to adapt to the environment independently. However, embodied drones can sense the environment in real time and respond to sudden changes. This ability makes them very useful in dangerous places like earthquake zones, forest fires, or floods. Recent studies show that embodied drones can perform many complex tasks. For instance, with the help of language models, they can understand and follow human voice instructions, helping drones quickly change their actions and enhancing their responses in emergency situations, such as “search near the collapsed bridge” [114], [239], [240], [241], [242]. Secondly, some work use world models to simulate dangerous environments, which helps them avoid danger and plan a safer path [243], [244]. Other studies explore how multiple drones can work together to find survivors and map damaged areas [197], [245].

However, despite these advancements, current approaches remain limited in handling long-horizon reasoning and autonomous decision-making under uncertainty. As illustrated in Fig. 7, jointing WMs and MLLMs has emerged as a key strategy for further enhancing UAV autonomy. The WM maintain a continuously evolving spatiotemporal representation of the environment, supporting planning and risk prediction even in GPS-denied conditions. The MLLM grounds commands into structured subtasks based on the UAV’s belief state. This coordination improves generalization, long-horizon reasoning, and high-level autonomy in mission-critical conditions.

C. Industrial Robots

Embodied AI is changing the way robots work in factories. With embodied AI, industrial robots can make

The purpose of autonomous embodied AI is to enable agents to operate independently for a long time in a dynamic and open environment.

smarter decisions based on their surroundings. Traditional industrial robots are usually fixed in one place. They use special sensors and tools and are required to complete tasks with very high accuracy. Because of this, they are better at doing jobs that need the same movements again and again.

However, with embodied AI, these robots can do more than repeat actions. By combining MLLMs and WMs, industrial robots can adjust how hard they hold fragile objects, or find a new path when they meet an obstacle. This has already been used in real life. For example, robots in Tesla's factory can find and fix parts that are not lined up, without help from people. At JD, robots [246], [247] use different sensors to sort packages by size and address. In Tmall's warehouse [248], robots use thermal cameras, LiDAR, and RGB sensors to check for problems in the inventory and send alerts when something is wrong. These examples show that embodied AI is helping robots become more flexible, reliable, and smart in factories.

D. Other Applications

In addition to its use in homes, healthcare, and rescue missions, embodied AI is also being applied in educational, virtual, and space environments. In smart manufacturing, it supports robots that can work together with humans, perform accurate assembly tasks, and adapt their actions based on changes in the workspace or human behavior. With the help of visual and touch feedback, these robots can safely handle fragile items [249], [250]. In education, embodied AI is used in social robots that adjust their speech, gaze, and gestures according to the student's focus and emotions. This helps create a more personalized learning experience and builds long-term trust between students and robots [251], [252]. In virtual environments, embodied agents learn to move, interact with objects, and complete tasks that require several steps. They also develop memory over time to improve their performance [253]. In space exploration, where conditions are unknown and communication with Earth is delayed, embodied AI allows robots to make decisions on their own and adapt to new surroundings [254]. These examples show that embodied AI is becoming more flexible and useful across many fields, helping machines see, act, and learn in both real and virtual worlds.

VII. Future Directions

As embodied AI moves from simulation to real-world deployment, future research must prioritize the development of unified and reliable systems across several core domains. Key directions include autonomous embodied AI, embodied AI hardware, swarm embodied AI, and evaluation benchmark.

A. Autonomous Embodied AI

The purpose of autonomous embodied AI is to enable agents to operate independently for a long time in a dynamic and open environment. Future research is expected to develop along several key directions. First, adaptive perception can give the system the ability to autonomously select input data, which can be achieved by dynamically choosing and integrating information from different sensory modalities. Second, Building on this foundation, building environmental awareness is essential. Environmental awareness helps agents quickly adapt to changes, predict the consequences of their actions, and transfer their behavior to new environments. It requires memory architectures that can capture spatiotemporal patterns and model causal relationships. Third, future systems should combine MLLMs with real-time physical interaction, which allows agents to bridge high-level language instructions with low-level control, and accurately model the real physical world.

B. Embodied AI Hardware

Future research in embodied AI hardware is expected to advance in the following four directions. First, hardware-aware model compression will continue to integrate techniques such as quantization and pruning with hardware performance metrics, enabling precise control over the trade-off between model accuracy and deployment efficiency. Second, graph-level compilation optimization will play a key role in bridging the gap between high-level embodied models and low-level hardware execution, which will focus on more effective operator fusion, scheduling strategies, and memory access efficiency to reduce execution overhead. Third, domain-specific accelerators will be increasingly tailored to the computational characteristics of embodied tasks. Reconfigurable architectures such as FPGA and CGRA offer flexibility and adaptability, while ASIC-based designs provide high efficiency and performance. Fourth, hardware-software co-design will become essential for eliminating mismatches between algorithm behavior and hardware architecture.

Joint optimization of model structures and hardware architecture will be critical to achieving real-time, energy-efficient execution in embodied systems.

C. Swarm Embodied AI

Swarm embodied AI refers to the collaborative perception and decision-making of multiple agents. Refers to the collaborative perception and decision-making of multiple agents. Because multiple agents can exhibit stronger capabilities when cooperating than a single agent, this kind of “collective intelligence” has aroused the interest of many researchers and is also regarded as an important step for agents to approach humans. First of all, to enable multiple agents to cooperate smoothly, it is necessary to develop collaborative WMs. This model can establish a shared and dynamic environmental representation based on the observations of each agent, forming the basis of collective understanding. Secondly, multi-agent representation learning is very important. It can help the agent understand its own state and also comprehend the situations of other agents. This is the basis for communication and cooperation among agents. In addition, modeling social behavior among agents is also crucial. Role allocation and group decision-making can be better achieved through behavioral modeling. Finally, to seamlessly integrate into real-world applications, it is also important to design natural human-swarm interaction interfaces. It may include multimodal language foundations and get-based control methods, making it easier for humans to direct and guide the entire agent group.

D. Explainability and Trustworthiness Embodied AI

Explainability and trustworthiness represent a critical frontier for Embodied AI, essential for its safe, ethical, and widespread real-world deployment as agents increasingly interact physically with humans and dynamic environments. Future research must address several key challenges: Firstly, designing benchmarks that provide real-time, human-understandable justifications for agent actions, particularly during unexpected situations or failures, is crucial for user trust and debugging. Secondly, establishing robust mechanisms to ensure agents adhere to ethical principles and human values during autonomous decision-making, especially in morally ambiguous scenarios common in rescue or healthcare applications, requires significant advancement. Thirdly, creating verifiable safety guarantees and certification standards for agents operating in unstructured physical settings, mitigating risks associated with unpredictable interactions, remains an open problem. Finally, enhancing robustness against adversarial attacks, sensor noise, and distribution shifts, ensuring reliable performance

despite uncertainties inherent in the real world, is fundamental for trustworthy operation. Addressing these multifaceted research problems in explainability and trustworthiness is paramount, as progress in this direction will unlock the full potential of Embodied AI by fostering user confidence, enabling responsible innovation, and facilitating regulatory acceptance.

E. Other Directions

Several new directions may influence the future development of embodied AI. One important direction is lifelong learning. Agents need to continuously learn new skills without forgetting what they have already learned. Only in this way can they adapt to the dynamic environment and maintain the accuracy of the previously completed tasks. Another key direction is human-in-the-loop learning. Human feedback is very important supervisory information. A small amount of feedback can significantly improve the performance of an agent and make it more human-like. To achieve this goal, we need better methods to enable agents to understand human goals and preferences. Finally, as agents become more autonomous, moral decision-making becomes increasingly important. Future systems should learn to carefully identify moral hazard and follow human values. This will help ensure that the embedded artificial intelligence is both safe and reliable.

Acknowledgment

The authors thank Dr. Ren Wang for his contribution to Section II. They also would like to thank Dr. Yu-Wei Zhan for her contributions to Sections VI and VII.



Tongtong Feng received the Ph.D. degree in computer science and technology from the Beijing University of Posts and Telecommunications. He is currently a Post-Doctoral Researcher with the Department of Computer Science and Technology, Tsinghua University.

He has published over 15 high-quality research papers in top journals and conferences, including IEEE TRANSACTIONS ON MULTIMEDIA, ESWA, ACM Multimedia, and AAAI. His research interests include embodied AI, world model, and multimedia intelligence. He received the Best Paper Nomination of ACM Multimedia in 2024.



Xin Wang (Member, IEEE) received the B.E. and Ph.D. degrees in computer science and technology from Zhejiang University, China, and the Ph.D. degree in computing science from Simon Fraser University, Canada. He is

currently an Associate Professor with the Department of Computer Science and Technology, Tsinghua University. He has published over 150 high-quality research papers in top journals and conferences, including IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE, IEEE TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING, ACM TOIS, ICML, NeurIPS, ACM KDD, ACM Web Conference, ACM SIGIR, and ACM Multimedia, and has won three best paper awards. His research interests include multimedia intelligence, machine learning, and its applications in multimedia big data analysis. He was a recipient of the 2020 ACM China Rising Star Award, the 2022 IEEE TCMC Rising Star Award, and the 2023 DAMO Academy Young Fellow.



Yu-Gang Jiang (Fellow, IEEE) received the Ph.D. degree in computer science from the City University of Hong Kong in 2009. He was a Post-Doctoral Research Scientist at Columbia University, New York, NY, USA, from 2009 to 2011. He is currently a Distinguished Professor of computer science at Fudan University, Shanghai, China. His research interests include multimedia, computer vision, embodied AI, and trustworthy AI. His research has led to the development of innovative AI tools that have been used in many practical applications, like defect detection for high-speed railway infrastructures. His open-source video analysis toolkits and datasets, such as CU-VIR-EO374, CCV, THUMOS, FCVID, and WildDeepfake, have been widely used in both academia and industry. He also serves as the Chair for the ACM Shanghai Chapter and an Associate Editor for several international journals. For contributions to large-scale and trustworthy video analysis, he was elected a fellow of IAPR and CCF.



Wenwu Zhu (Fellow, IEEE) received the Ph.D. degree from New York University in 1996. He was a Senior Researcher and the Research Manager at Microsoft Research Asia. He was the Chief Scientist and the Director of Intel Research China from 2004 to 2008. He was a Member of Technical Staff at Bell Labs, Murray Hill, NJ, USA, from 1996 to 1999. He is currently a Professor with the Department of Computer Science and Technology, Tsinghua University, and the Vice Dean of the Beijing National Research Center for Information Science and Technology. He has published over 400 refereed articles and is an inventor of over 100 patents. His research interests include graph machine learning, curriculum learning, data-driven multimedia, and big data. He is an AAAS Fellow, ACM Fellow, SPIE Fellow, and a member of Academia

Europaea. He received ten Best Paper Awards, including ACM Multimedia 2012 and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY in 2001 and 2019. He serves as the General Co-Chair for ACM Multimedia 2018 and ACM CIKM 2019. He served as the Editor-in-Chief for IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS FOR VIDEO TECHNOLOGY from 2024 to 2025 and IEEE TRANSACTIONS ON MULTIMEDIA from 2017 to 2019, and the Chair for the Steering Committee of IEEE TRANSACTIONS ON MULTIMEDIA from 2020 to 2022.

References

- [1] A. M. Turing, "Computing machinery and intelligence," *Mind*, vol. 59, no. 236, pp. 433–460, 1950.
- [2] R. Pfeifer and J. Bongard, *How the Body Shapes the Way We Think: A New View of Intelligence*. Cambridge, MA, USA: MIT Press, 2006.
- [3] A. Clark, *Being There: Putting Brain, Body, and World Together Again*. Cambridge, MA, USA: MIT Press, 1998.
- [4] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," in *Proc. Int. Conf. Learn. Represent.*, 2015, pp. 1–14.
- [5] T. Chen et al., "TVM: An automated end-to-end optimizing compiler for deep learning," in *Proc. 13th USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2018, pp. 578–594.
- [6] N. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," in *Proc. 44th Annu. Int. Symp. Comput. Architect.*, 2017, pp. 1–12.
- [7] J. Engel, T. Schöps, and D. Cremers, "LSD-SLAM: Large-scale direct monocular SLAM," in *Proc. Eur. Conf. Comput. Vis.*, Zurich, Switzerland, Sep. 2014, pp. 834–849.
- [8] R. Mur-Artal, J. M. M. Montiel, and J. D. Tardos, "ORB-SLAM: A versatile and accurate monocular SLAM system," *IEEE Trans. Robot.*, vol. 31, no. 5, pp. 1147–1163, Oct. 2015.
- [9] A. Macario Barros et al., "A comprehensive survey of visual SLAM algorithms," *Robotics*, vol. 11, no. 1, p. 24, Feb. 2022.
- [10] Z. Wu et al., "Embodied task planning with large language models," 2023, *arXiv:2307.01848*.
- [11] S. Wang et al., "OmniDrive: A holistic vision-language dataset for autonomous driving with counterfactual reasoning," 2024, *arXiv:2405.01533*.
- [12] D. Jayaraman and K. Grauman, "Learning to look around: Intelligently exploring unseen environments for unknown tasks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1238–1247.
- [13] B. Y. Lin et al., "On grounded planning for embodied tasks with language models," in *Proc. AAAI Conf. Artif. Intell.*, 2023, vol. 37, no. 11, pp. 13192–13200.
- [14] D. Driess et al., "Palm-e: An embodied multimodal language model," in *Proc. Int. Conf. Mach. Learn.*, 2023, pp. 8469–8488.
- [15] W. Jiang et al., "Multimodal LLM guided exploration and active mapping using Fisher information," 2024, *arXiv:2410.17422*.
- [16] Z. Wang et al., "ActiView: Evaluating active perception ability for multimodal large language models," 2024, *arXiv:2410.04659*.
- [17] Y. Qin et al., "MP5: A multi-modal open-ended embodied system in mincraft via active perception," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 16307–16316.
- [18] R. Yang et al., "EmbodiedBench: Comprehensive benchmarking multi-modal large language models for vision-driven embodied agents," 2025, *arXiv:2502.09560*.
- [19] A. Brohan et al., "RT-1: Robotics transformer for real-world control at scale," 2022, *arXiv:2212.06817*.
- [20] J.-B. Alayrac et al., "Flamingo: A visual language model for few-shot learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 23716–23736.
- [21] L. Floridi and M. Chiriatti, "GPT-3: Its nature, scope, limits, and consequences," *Minds Mach.*, vol. 30, no. 4, pp. 681–694, Dec. 2020.
- [22] M. Ahn et al., "Do as I can, not as I say: Grounding language in robotic affordances," 2022, *arXiv:2204.01691*.
- [23] I. Singh et al., "ProgPrompt: Program generation for situated robot task planning using large language models," *Auto. Robots*, vol. 47, no. 8, pp. 999–1012, Dec. 2023.

- [24] A. Brohan et al., "RT-2: Vision-language-action models transfer web knowledge to robotic control," 2023, *arXiv:2307.15818*.
- [25] J. Ding et al., "Understanding world or predicting future? A comprehensive survey of world models," 2024, *arXiv:2411.14499*.
- [26] D. Ha and J. Schmidhuber, "World models," 2018, *arXiv:1803.10122*.
- [27] Y. LeCun, "A path towards autonomous machine intelligence version 0.9. 2, 2022-06-27," *Open Rev.*, vol. 62, no. 1, pp. 1–62, 2022.
- [28] X. Li et al., "ManipLLM: Embodied multimodal large language model for object-centric robotic manipulation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 18061–18070.
- [29] J. Duan et al., "A survey of embodied AI: From simulators to research tasks," *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 6, no. 2, pp. 230–244, Apr. 2022.
- [30] T. Brophy et al., "A review of the impact of rain on camera-based perception in automated driving systems," *IEEE Access*, vol. 11, pp. 67040–67057, 2023.
- [31] K. Black et al., " π 0: A vision-language-action flow model for general robot control," 2024, *arXiv:2410.24164*.
- [32] W. Huang et al., "ReKep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation," 2024, *arXiv:2409.01652*.
- [33] Q. Zhao et al., "CoT-VLA: Visual chain-of-thought reasoning for vision-language-action models," 2025, *arXiv:2503.22020*.
- [34] Y. Liu, J. Smith, and Q. Chen, "Task alignment in embodied AI: Bridging semantics and physical constraints," *Trans. Robot.*, vol. 40, pp. 1–15, Jan. 2024.
- [35] D. Driess, Y. Huang, and A. E. A. Zeng, "Vision-language-action models: Unified frameworks for embodied execution," in *Proc. Conf. Robot Learn. (CoRL)*, 2023, pp. 1234–1245.
- [36] J. Bruce et al., "Genie: Generative interactive environments," 2024, *arXiv:2402.15391*.
- [37] C. Chen et al., "TransDreamer: Reinforcement learning with transformer world models," 2022, *arXiv:2202.09481*.
- [38] J. Robine et al., "Transformer-based world models are happy with 100k interactions," in *Proc. ICLR*, 2023, pp. 1–21.
- [39] X. Wang et al., "WorldDreamer: Towards general world models for video generation via predicting masked tokens," 2024, *arXiv:2401.09985*.
- [40] W. Wu et al., "Paragraph-to-image generation with information-enriched diffusion model," 2023, *arXiv:2311.14284*.
- [41] D. Hafner et al., "Dream to control: Learning behaviors by latent imagination," in *Proc. ICLR*, 2020, pp. 1–20.
- [42] D. Hafner et al., "Mastering Atari with discrete world models," in *Proc. ICLR*, 2021, pp. 1–26.
- [43] M. Okada and T. Taniguchi, "DreamingV2: Reinforcement learning with discrete world models without reconstruction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2022, pp. 985–991.
- [44] P. Wu et al., "Daydreamer: World models for physical robot learning," in *Proc. Conf. Robot Learn.*, 2023, pp. 2226–2240.
- [45] A. Clark, *Whatever Next? Predictive Brains and Embodied Cognition*. Cambridge, MA, USA: MIT Press, 2013.
- [46] T. Feng et al., "EvoAgent: Agent autonomous evolution with continual world model for long-horizon tasks," 2025, *arXiv:2502.05907*.
- [47] G. Lakoff and M. Johnson, *Metaphors We Live*. Chicago, IL, USA: Univ. Chicago Press, 1980.
- [48] S. Harnad, "The symbol grounding problem," *Phys. D, Nonlinear Phenomena*, vol. 42, nos. 1–3, pp. 335–346, Jun. 1990.
- [49] R. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robot. Autom.*, vol. RA-2, no. 1, pp. 14–23, Jan. 1986.
- [50] R. A. Brooks, "Intelligence without representation," *Artif. Intell.*, vol. 47, nos. 1–3, pp. 139–159, Jan. 1991.
- [51] R. A. Brooks et al., "The cog project: Building a humanoid robot," in *Proc. Int. Workshop Comput. Metaphors, Analogy, Agents*. Cham, Switzerland: Springer, 1998, pp. 52–87.
- [52] R. Chrisley, "Embodied artificial intelligence," *Artif. Intell.*, vol. 149, no. 1, pp. 131–150, Sep. 2003.
- [53] S. Levine et al., "End-to-end training of deep visuomotor policies," *J. Mach. Learn. Res.*, vol. 17, no. 39, pp. 1–40, 2016.
- [54] V. Mnih et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [55] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 25, 2012, pp. 1–11.
- [56] I. J. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 27, 2014, pp. 52–87.
- [57] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [58] A. Dosovitskiy et al., "An image is worth 16x16 words: Transformers for image recognition at scale," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–22.
- [59] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 33, 2020, pp. 6840–6851.
- [60] K. He et al., "Masked autoencoders are scalable vision learners," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 16000–16009.
- [61] A. Kirillov et al., "Segment anything," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Oct. 2023, pp. 4015–4026.
- [62] A. Vaswani, "Attention is all you need," in *Proc. NeurIPS*, 2017, pp. 1–11.
- [63] J. Devlin, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol. (Long Short Papers)*, vol. 1, 2019, pp. 4171–4186.
- [64] C. Raffel et al., "Exploring the limits of transfer learning with a unified text-to-text transformer," *J. Mach. Learn. Res.*, vol. 21, no. 140, pp. 1–67, 2020.
- [65] J. Achiam et al., "GPT-4 technical report," 2023, *arXiv:2303.08774*.
- [66] W.-L. Chiang et al. (2023). *Vicuna: An Open-Source Chatbot Impressing GPT-4 With 90% ChatGPT Quality*. Accessed: Apr. 14, 2023. [Online]. Available: <https://vicuna.lmsys.org>
- [67] H. Touvron et al., "LLaMA: Open and efficient foundation language models," 2023, *arXiv:2302.13971*.
- [68] V. Mnih et al., "Playing Atari with deep reinforcement learning," 2013, *arXiv:1312.5602*.
- [69] D. Silver et al., "Mastering the game of go with deep neural networks and tree search," *Nature*, vol. 529, no. 7587, pp. 484–489, Jan. 2016.
- [70] J. Schulman et al., "Proximal policy optimization algorithms," 2017, *arXiv:1707.06347*.
- [71] T. Haarnoja et al., "Soft actor-critic: Offpolicy maximum entropy deep reinforcement learning with a stochastic actor," in *Proc. Int. Conf. Mach. Learn.*, vol. 80, 2018, pp. 1861–1870.
- [72] L. Ouyang et al., "Training language models to follow instructions with human feedback," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2022, pp. 27730–27744.
- [73] Z. Shao et al., "DeepSeekMath: Pushing the limits of mathematical reasoning in open language models," 2024, *arXiv:2402.03300*.
- [74] J. Bai et al., "Qwen technical report," 2023, *arXiv:2309.16609*.
- [75] G. Team et al., "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," 2024, *arXiv:2403.05530*.
- [76] A. Hurst et al., "GPT-4o system card," 2024, *arXiv:2410.21276*.
- [77] D. Guo et al., "DeepSeek-r1: Incentivizing reasoning capability in LLMs via reinforcement learning," 2025, *arXiv:2501.12948*.
- [78] D. Hafner et al., "Learning latent dynamics for planning from pixels," in *Proc. Int. Conf. Mach. Learn.*, May 2019, pp. 2555–2565.
- [79] D. Hafner, J. Pasukonis, J. Ba, and T. Lillicrap, "Mastering diverse domains through world models," 2023, *arXiv:2301.04104*.
- [80] Z. Zhu et al., "Is sora a world simulator? A comprehensive survey on general world models and beyond," 2024, *arXiv:2405.03520*.
- [81] D. S. Chaplot et al., "SEAL: Self-supervised embodied active learning using exploration and 3d consistency," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 13086–13098.
- [82] M. F. Ahmed et al., "Active slam: A review on last decade," *Sensors*, vol. 23, no. 19, p. 8097, 2023.
- [83] M. R. U. Saputra, A. Markham, and N. Trigoni, "Visual SLAM and structure from motion in dynamic environments: A survey," *ACM Comput. Surv.*, vol. 51, no. 2, pp. 1–36, Mar. 2019.
- [84] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: Part I," *IEEE Robot. Autom. Mag.*, vol. 13, no. 2, pp. 99–110, Jun. 2006.
- [85] T. Bailey and H. Durrant-Whyte, "Simultaneous localization and mapping (SLAM): Part II," *IEEE Robot. Autom. Mag.*, vol. 13, no. 3, pp. 108–117, Sep. 2006.
- [86] Y. Wang et al., "A survey of visual SLAM in dynamic environment: The evolution from geometric to semantic approaches," *IEEE Trans. Instrum. Meas.*, vol. 73, pp. 1–21, 2024.
- [87] P. F. Alcantarilla et al., "On combining visual SLAM and dense scene flow to increase the robustness of localization and mapping in dynamic environments," in *Proc. IEEE Int. Conf. Robot. Autom.*, May 2012, pp. 1290–1297.

- [88] D.-H. Kim and J.-H. Kim, "Effective background model-based RGB-D dense visual odometry in a dynamic environment," *IEEE Trans. Robot.*, vol. 32, no. 6, pp. 1565–1573, Dec. 2016.
- [89] D. Zou and P. Tan, "CoSLAM: Collaborative visual SLAM in dynamic environments," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 2, pp. 354–366, Feb. 2013.
- [90] W. Dai, Y. Zhang, P. Li, Z. Fang, and S. Scherer, "RGB-D SLAM in dynamic environments using point correlations," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 1, pp. 373–389, Jan. 2022.
- [91] Z.-J. Du et al., "Accurate dynamic SLAM using CRF-based long-term consistency," *IEEE Trans. Vis. Comput. Graphics*, vol. 28, no. 4, pp. 1745–1757, Apr. 2022.
- [92] R. F. Salas-Moreno et al., "SLAM++: Simultaneous localisation and mapping at the level of objects," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 1352–1359.
- [93] C. Yu et al., "DSSLAM: A semantic visual slam towards dynamic environments," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2018, pp. 1168–1174.
- [94] M. Gonzalez, E. Marchand, A. Kacete, and J. Royan, "TwistSLAM: Constrained SLAM in dynamic environment," *IEEE Robot. Autom. Lett.*, vol. 7, no. 3, pp. 6846–6853, Jul. 2022.
- [95] C. Yan et al., "GS-SLAM: Dense visual SLAM with 3D Gaussian splatting," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 19595–19604.
- [96] M. A. Bautista et al., "GAUDI: A neural architect for immersive 3D scene generation," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 25102–25116.
- [97] R. Chen et al., "Towards label-free scene understanding by vision foundation models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, p. 75.
- [98] S. Peng et al., "OpenScene: 3D scene understanding with open vocabularies," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 815–824.
- [99] Y. Man et al., "Lexicon3D: Probing visual foundation models for complex 3D scene understanding," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, 2024, pp. 76819–76847.
- [100] G. Gao et al., "GraphDreamer: Compositional 3D scene synthesis from scene graphs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 21295–21304.
- [101] H. Zhou et al., "HUGS: Holistic urban 3D scene understanding via Gaussian splatting," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 21336–21345.
- [102] J. Yang et al., "RegionPLC: Regional point-language contrastive learning for open-world 3D scene understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 19823–19832.
- [103] P. Shyam, W. Jaśkowski, and F. Gomez, "Model-based active exploration," in *Proc. 36th Int. Conf. Mach. Learn.*, 2019, pp. 5779–5788.
- [104] D. S. Chaplot et al., "Learning to explore using active neural slam," in *Proc. 8th Int. Conf. Learn. Represent.*, 2020, pp. 1–18.
- [105] H. Liu and P. Abbeel, "Behavior from the void: Unsupervised active pre-training," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 18459–18473.
- [106] M. Xu et al., "Active reasoning in an open-world environment," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, pp. 11716–11736.
- [107] A. Russo and A. Proutiere, "Model-free active exploration in reinforcement learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, pp. 54740–54753.
- [108] A. Somayazulu et al., "ActiveRIR: Active audio-visual exploration for acoustic environment modeling," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2024, pp. 13830–13836.
- [109] L. Zhang, G. Yang, and B. C. Stadie, "World model as a graph: Learning latent landmarks for planning," in *Proc. Int. Conf. Mach. Learn.*, 2021, pp. 12611–12620.
- [110] C. H. Song et al., "LLM-planner: Few-shot grounded planning for embodied agents with large language models," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2023, pp. 2986–2997.
- [111] X. Zhou et al., "EGO-planner: An ESDF-free gradient-based local planner for quadrotors," *IEEE Robot. Autom. Lett.*, vol. 6, no. 2, pp. 478–485, Apr. 2021.
- [112] S. Qiao et al., "AutoAct: Automatic agent learning from scratch for QA via self-planning," 2024, [arXiv:2401.05268](https://arxiv.org/abs/2401.05268).
- [113] L. Yang et al., "Mastering text-to-image diffusion: Recaptioning, planning, and generating with multimodal llms," in *Proc. 41st Int. Conf. Mach. Learn.*, 2024, pp. 1–18.
- [114] D. An et al., "ETPNav: Evolving topological planning for vision-language navigation in continuous environments," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 47, no. 7, pp. 5130–5145, Jul. 2025.
- [115] N. Shinn et al., "Reflexion: Language agents with verbal reinforcement learning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 8634–8652.
- [116] Z. Liu, A. Bahety, and S. Song, "REFLECT: Summarizing robot experiences for failure explanation and correction," 2023, [arXiv:2306.15724](https://arxiv.org/abs/2306.15724).
- [117] Z. Yang et al., "RILA: Reflective and imaginative language agent for zero-shot semantic audio-visual navigation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 16251–16261.
- [118] Z. Li, Y. Xie, R. Shao, G. Chen, D. Jiang, and L. Nie, "Optimus-1: Hybrid multimodal memory empowered agents excel in long-horizon tasks," 2024, [arXiv:2408.03615](https://arxiv.org/abs/2408.03615).
- [119] P. Yuan et al., "REMAC: Self-reflective and self-evolving multi-agent collaboration for long-horizon robot manipulation," 2025, [arXiv:2503.22122](https://arxiv.org/abs/2503.22122).
- [120] A. Brohan et al., "Do as I can, not as I say: Grounding language in robotic affordances," in *Proc. Conf. Robot Learn.*, 2023, pp. 287–318.
- [121] S. Reed et al., "A generalist agent," 2022, [arXiv:2205.06175](https://arxiv.org/abs/2205.06175).
- [122] Y. Mu et al., "EmbodiedGPT: Vision-language pre-training via embodied chain of thought," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 25081–25094.
- [123] Z. Peng et al., "Kosmos-2: Grounding multimodal large language models to the world," 2023, [arXiv:2306.14824](https://arxiv.org/abs/2306.14824).
- [124] Y. Hong et al., "MultiPLY: A multisensory object-centric embodied large language model in 3D world," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 26396–26406.
- [125] L. Fan et al., "MineDojo: Building open-ended embodied agents with internet-scale knowledge," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 18343–18362.
- [126] M. Jin Kim et al., "OpenVLA: An open-source vision-language-action model," 2024, [arXiv:2406.09246](https://arxiv.org/abs/2406.09246).
- [127] W. Hong et al., "CogAgent: A visual language model for GUI agents," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 14281–14290.
- [128] D. Ghosh et al., "Octo: An open-source generalist robot policy," 2024, [arXiv:2405.12213](https://arxiv.org/abs/2405.12213).
- [129] R. Doshi et al., "Scaling cross-embodied learning: One policy for manipulation, navigation, locomotion and aviation," 2024, [arXiv:2408.11812](https://arxiv.org/abs/2408.11812).
- [130] L. Wang et al., "Scaling proprioceptive-visual learning with heterogeneous pre-trained transformers," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, 2024, pp. 124420–124450.
- [131] J. Ho and S. Ermon, "Generative adversarial imitation learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, 2016, pp. 1–9.
- [132] N. Baram et al., "End-to-end differentiable adversarial imitation learning," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 390–399.
- [133] S. Suo et al., "TrafficSim: Learning to simulate realistic multi-agent behaviors," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 10395–10404.
- [134] Q. Zhang et al., "TrajGen: Generating realistic and diverse trajectories with reactive and feasible agent behaviors for autonomous driving," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 12, pp. 24474–24487, Dec. 2022.
- [135] C. Li et al., "Behavior-1k: A benchmark for embodied AI with 1,000 everyday activities and realistic simulation," in *Proc. Conf. Robot Learn.*, 2023, pp. 80–93.
- [136] J. Lu et al., "AgentLens: Visual analysis for agent behaviors in LLM-based autonomous systems," *IEEE Trans. Vis. Comput. Graphics*, vol. 31, no. 8, pp. 4182–4197, Aug. 2025.
- [137] B. Chen et al., "Embodied contrastive learning with geometric consistency and behavioral awareness for object navigation," in *Proc. 32nd ACM Int. Conf. Multimedia*, Oct. 2024, pp. 4776–4785.
- [138] T. Rashid et al., "Monotonic value function factorisation for deep multiagent reinforcement learning," *J. Mach. Learn. Res.*, vol. 21, no. 178, pp. 1–51, 2020.
- [139] K. Son et al., "QTRAN: Learning to factorize with transformation for cooperative multi-agent reinforcement learning," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 5887–5896.
- [140] J. Wang et al., "QPLEX: Duplex dueling multi-agent Q-learning," 2020, [arXiv:2008.01062](https://arxiv.org/abs/2008.01062).
- [141] M. Wen et al., "Multi-agent reinforcement learning is a sequence modeling problem," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 35, 2022, pp. 16509–16521.

- [142] H. Zhang et al., "Building cooperative embodied agents modularly with large language models," 2023, *arXiv:2307.02485*.
- [143] W. Chen et al., "AgentVerse: Facilitating multi-agent collaboration and exploring emergent behaviors," 2023, *arXiv:2308.10848*.
- [144] S. Hong et al., "MetaGPT: Meta programming for a multi-agent collaborative framework," 2023, *arXiv:2308.00352*.
- [145] H. Zhang et al., "COMBO: Compositional world models for embodied multi-agent cooperation," 2024, *arXiv:2404.10775*.
- [146] A. Clark, "An embodied cognitive science?" *Trends Cognitive Sci.*, vol. 3, no. 9, pp. 345–351, 1999.
- [147] R. W. Gibbs Jr., *Embodiment and Cognitive Science*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [148] T. Ziemke, "What's that thing called embodiment?" in *Proc. 25th Annu. Cognitive Sci. Soc.* Milton Park, U.K.: Psychology Press, 2013, pp. 1305–1310.
- [149] X. Liu, H. Palacios, and C. Muise, "Egocentric planning for scalable embodied task achievement," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, pp. 54586–54613.
- [150] K. C. Stocking, A. Gopnik, and C. Tomlin, "From robot learning to robot understanding: Leveraging causal graphical models for robotics," in *Proc. Conf. Robot Learn.*, 2022, pp. 1776–1781.
- [151] P. Sermanet et al., "RoboVQA: Multimodal long-horizon reasoning for robotics," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2024, pp. 645–652.
- [152] J. Zhang et al., "FLTRNN: Faithful long-horizon task planning for robotics with large language models," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2024, pp. 6680–6686.
- [153] S. Nayak et al., "Long-horizon planning for multi-agent robots in partially observable environments," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, 2024, pp. 67929–67967.
- [154] X. Zhang et al., "LaMMA-P: Generalizable multi-agent long-horizon task allocation and planning with LM-driven PDDL planner," 2024, *arXiv:2409.20560*.
- [155] A. Padmakumar et al., "TEACH: Task-driven embodied agents that chat," in *Proc. AAAI Conf. Artif. Intell.*, 2022, vol. 36, no. 2, pp. 2017–2025.
- [156] Z. Zhang et al., "A survey on the memory mechanism of large language model based agents," 2024, *arXiv:2404.13501*.
- [157] B. Bakker, "Reinforcement learning with long short-term memory," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 14, 2001, pp. 1–8.
- [158] D. Ramani, "A short survey on memory based reinforcement learning," 2019, *arXiv:1904.06736*.
- [159] G. Zhu et al., "Episodic reinforcement learning with associative memory," in *Proc. Int. Conf. Learn. Represent.*, 2020, pp. 1–15.
- [160] A. Khan et al., "Memory augmented control networks," 2017, *arXiv:1709.05706*.
- [161] A. Pritzel et al., "Neural episodic control," in *Proc. Int. Conf. Mach. Learn.*, vol. 70. Sydney, NSW, Australia, Aug. 2017, pp. 2827–2836.
- [162] Z. Xu et al., "A survey on robotics with foundation models: Toward embodied AI," 2024, *arXiv:2402.02385*.
- [163] L. Ren et al., "Embodied intelligence toward future smart manufacturing in the era of AI foundation model," *IEEE/ASME Trans. Mechatronics*, vol. 30, no. 4, pp. 2632–2642, Aug. 2025.
- [164] R. Firoozi et al., "Foundation models in robotics: Applications, challenges, and the future," *Int. J. Robot. Res.*, vol. 44, no. 5, pp. 701–739, Apr. 2025.
- [165] K. Li et al., "MuEP: A multimodal benchmark for embodied planning with foundation models," in *Proc. 33rd Int. Joint Conf. Artif. Intell.*, Aug. 2024, pp. 129–138.
- [166] R. Dang et al., "ECBench: Can multi-modal foundation models understand the egocentric world? A holistic embodied cognition benchmark," 2025, *arXiv:2501.05031*.
- [167] M. Zhang et al., "MFE-ETP: A comprehensive evaluation benchmark for multi-modal foundation models on embodied task planning," 2024, *arXiv:2407.05047*.
- [168] F. Sun, R. Chen, T. Ji, Y. Luo, H. Zhou, and H. Liu, "A comprehensive survey on embodied intelligence: Advancements, challenges, and future perspectives," *CAAI Artif. Intell. Res.*, vol. 3, Dec. 2024, Art. no. 9150042.
- [169] L. Jin and L. Jia, "Embodied world models emerge from navigational task in open-ended environments," 2025, *arXiv:2504.11419*.
- [170] A. Billard and D. Kragic, "Trends and challenges in robot manipulation," *Science*, vol. 364, no. 6446, 2019, Art. no. eaat8414.
- [171] A. Ajoudani, A. M. Zanchettin, S. Ivaldi, A. Albu-Schäffer, K. Kotsuge, and O. Khatib, "Progress and prospects of the human-robot collaboration," *Auto. Robots*, vol. 42, no. 5, pp. 957–975, Jun. 2018.
- [172] R. Kelly, "A tuning procedure for stable PID control of robot manipulators," *Robotica*, vol. 13, no. 2, pp. 141–148, Mar. 1995.
- [173] P. Rocco, "Stability of PID control for industrial robot arms," *IEEE Trans. Robot. Autom.*, vol. 12, no. 4, pp. 606–614, Jun. 1996.
- [174] B. Kiumarsi, K. G. Vamvoudakis, H. Modares, and F. L. Lewis, "Optimal and autonomous control using reinforcement learning: A survey," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 6, pp. 2042–2062, Jun. 2018.
- [175] B. Singh, R. Kumar, and V. P. Singh, "Reinforcement learning in robotic applications: A comprehensive survey," *Artif. Intell. Rev.*, vol. 55, no. 2, pp. 945–990, Feb. 2022.
- [176] Y. Ma, Z. Song, Y. Zhuang, J. Hao, and I. King, "A survey on vision-language-action models for embodied AI," 2024, *arXiv:2405.14093*.
- [177] B. Jiang et al., "Multi-modal and multi-agent systems meet rationality: A survey," in *Proc. ICML Workshop LLMs Cognition*, 2024, pp. 1–16.
- [178] Z. Zhu et al., "MADiff: Offline multi-agent learning with diffusion models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, 2024, pp. 4177–4206.
- [179] M. Ma and L. Cheng, "A human-robot collaboration controller utilizing confidence for disagreement adjustment," *IEEE Trans. Robot.*, vol. 40, pp. 2081–2097, 2024.
- [180] G. Xiao et al., "SmoothQuant: Accurate and efficient post-training quantization for large language models," in *Proc. Int. Conf. Mach. Learn.*, 2023, p. 38.
- [181] H. Wang, Z. Zhang, and S. Han, "SpAtten: Efficient sparse attention architecture with cascade token and head pruning," in *Proc. IEEE Int. Symp. High-Performance Comput. Archit. (HPCA)*, Feb. 2021, pp. 97–110.
- [182] K. Wang, Z. Liu, Y. Lin, J. Lin, and S. Han, "HAQ: Hardware-aware automated quantization with mixed precision," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8612–8620.
- [183] C. Lattner and V. Adve, "LLVM: A compilation framework for life-long program analysis & transformation," in *Proc. Int. Symp. Code Gener. Optim. (CGO)*, 2004, pp. 75–86.
- [184] Q. Sun et al., "GTuner: Tuning DNN computations on GPU via graph attention network," in *Proc. 59th ACM/IEEE Design Autom. Conf.*, Jul. 2022, pp. 1045–1050.
- [185] J. Cai et al., "Gemini: Mapping and architecture co-exploration for large-scale DNN chiplet accelerators," in *Proc. IEEE Int. Symp. High-Perform. Comput. Archit. (HPCA)*, Mar. 2024, pp. 156–171.
- [186] J. Liu et al., "FlightVGM: Efficient video generation model inference with online sparsification and hybrid precision on FPGAs," in *Proc. ACM/SIGDA Int. Symp. Field Program. Gate Arrays*, Feb. 2025, pp. 2–13.
- [187] J. Qin, T. Xia, C. Tan, J. Zhang, and S. Q. Zhang, "PICACHU: Plug-in CGRA handling upcoming nonlinear operations in LLMs," in *Proc. 30th ACM Int. Conf. Architectural Support Program. Lang. Operating Syst., Volume 2*, Mar. 2025, pp. 845–861.
- [188] H. Genc et al., "Gemmini: Enabling systematic deep-learning architecture evaluation via full-stack integration," in *Proc. 58th ACM/IEEE Design Autom. Conf. (DAC)*, Dec. 2021, pp. 769–774.
- [189] Y. Lin et al., "QServe: W4A8KV4 quantization and system co-design for efficient LLM serving," 2024, *arXiv:2405.04532*.
- [190] M. Wang et al., "JAQ: Joint efficient architecture design and low-bit quantization with hardware-software co-exploration," 2025, *arXiv:2501.05339*.
- [191] L. Zou et al., "BiE: Bi-exponent block floating-point for large language models quantization," in *Proc. 41st Int. Conf. Mach. Learn.*, 2024, pp. 1–19.
- [192] M. Savva et al., "Habitat: A platform for embodied AI research," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Mali, Oct. 2019, pp. 9338–9346.
- [193] T. Mu et al., "ManiSkill: Learning-from-demonstrations benchmark for generalizable manipulation skills," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–22.
- [194] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, Oct. 2012, pp. 5026–5033.
- [195] S. Shah, D. Dey, C. Lovett, and A. Kapoor, "AirSim: High-fidelity visual and physical simulation for autonomous vehicles," in *Field and Service Robotics*. Cham, Switzerland: Springer, 2018, pp. 621–635.
- [196] M. Jin Kim et al., "OpenVLA: An open-source vision-language-action model," 2024, *arXiv:2406.09246*.
- [197] Y. Zhang et al., "Towards efficient LLM grounding for embodied multi-agent collaboration," 2024, *arXiv:2405.14314*.
- [198] T. Wang et al., "EmbodiedScan: A holistic multi-modal 3D perception suite towards embodied AI," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 19757–19767.

- [199] X. Wang et al., "Drivedreamer: Towards real-world-drive world models for autonomous driving," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2024, pp. 55–72.
- [200] J. Thomason et al., "Language grounding with 3D objects," in *Proc. Conf. Robot Learn.*, 2022, pp. 1691–1701.
- [201] S. Wang et al., "LLM³: Large language model-based task and motion planning with motion failure reasoning," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2024, pp. 12086–12092.
- [202] J. Hong, R. Choi, and J. J. Leonard, "Semantic enhancement for object SLAM with heterogeneous multimodal large language model agents," 2024, *arXiv:2411.06752*.
- [203] J. Yang et al., "Magma: A foundation model for multimodal AI agents," 2025, *arXiv:2502.13130*.
- [204] J. Yue et al., "MLLM as retriever: Interactively learning multimodal retrieval for embodied agents," 2024, *arXiv:2410.03450*.
- [205] Z. Chen et al., "Intern VL: Scaling up vision foundation models and aligning for generic visual-linguistic tasks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2024, pp. 24185–24198.
- [206] J. Xu et al., "Qwen2.5-omni technical report," 2025, *arXiv:2503.20215*.
- [207] B. Zitkovich et al., "RT-2: Vision-language-action models transfer web knowledge to robotic control," in *Proc. Conf. Robot Learn.*, 2023, pp. 2165–2183.
- [208] X. Wang et al., "Executable code actions elicit better LLM agents," in *Proc. 41st Int. Conf. Mach. Learn.*, 2024, pp. 1–13.
- [209] D. Wu et al., "Generative multi-agent collaboration in embodied AI: A systematic review," 2025, *arXiv:2502.11518*.
- [210] W. Zhao et al., "VLAS: Vision-language-action model with speech instructions for customized robot manipulation," 2025, *arXiv:2502.13508*.
- [211] H. Liu et al., "Visual instruction tuning," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 34892–34916.
- [212] J. W. Forrester, "Counterintuitive behavior of social systems," *Theory Decision*, vol. 2, no. 2, pp. 109–140, 1971.
- [213] M. Assran et al., "Self-supervised learning from images with a joint-embedding predictive architecture," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2023, pp. 15619–15629.
- [214] R. Yang, P. Srivastava, and S. Mandt, "Diffusion probabilistic modeling for video generation," *Entropy*, vol. 25, no. 10, p. 1469, Oct. 2023.
- [215] J. Cho et al., "Sora as an AGI world model? A complete survey on text-to-video generation," 2024, *arXiv:2403.05131*.
- [216] D. Li et al., "WorldModelBench: Judging video generation models as world models," 2025, *arXiv:2502.20694*.
- [217] J. Xiang et al., "Language models meet world models: Embodied experiences enhance language models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, pp. 75392–75412.
- [218] P. Mazzaglia et al., "GenRL: Multimodal-foundation world models for generalization in embodied agents," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, 2024, pp. 27529–27555.
- [219] T. Gupta et al., "The essential role of causality in foundation world models for embodied AI," 2024, *arXiv:2402.06665*.
- [220] N. Hansen et al., "Hierarchical world models as visual whole-body humanoid controllers," 2024, *arXiv:2405.18418*.
- [221] E. Alonso et al., "Diffusion for world modeling: Visual details matter in Atari," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, pp. 58757–58791, 2024.
- [222] A. Bardes et al., "Revisiting feature prediction for learning visual representations from video," 2024, *arXiv:2404.08471*.
- [223] A. Bardes, J. Ponce, and Y. LeCun, "MC-JEPA: A joint-embedding predictive architecture for self-supervised learning of motion and content features," 2023, *arXiv:2307.12698*.
- [224] Z. Fei, M. Fan, and J. Huang, "A-JEPA: Joint-embedding predictive architecture can listen," 2023, *arXiv:2311.15830*.
- [225] A. Banino et al., "MEMO: A deep network for flexible combination of episodic memories," 2020, *arXiv:2001.10913*.
- [226] V. Micheli, E. Alonso, and F. Fleuret, "Transformers are sample-efficient world models," in *Proc. ICLR*, 2023, pp. 1–21.
- [227] M. Xu et al., "Spatial-temporal transformer networks for traffic flow forecasting," 2020, *arXiv:2001.02908*.
- [228] A. Radford et al., "Learning transferable visual models from natural language supervision," in *Proc. Int. Conf. Mach. Learn.*, vol. 139, 2021, pp. 8748–8763.
- [229] S. Huang et al., "Language is not all you need: Aligning perception with language models," in *Proc. Adv. Neural Inf. Process. Syst.*, 2023, pp. 72096–72109.
- [230] J. Liang et al., "Code as policies: Language model programs for embodied control," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2023, pp. 9493–9500.
- [231] H.-Y. Tung et al., "Physion++: Evaluating physical scene understanding that requires online inference of different physical properties," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, 2023, pp. 67048–67068.
- [232] X. Shi et al., "Are we ready for service robots? The OpenLORIS-scene datasets for lifelong SLAM," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 3139–3145.
- [233] R. Wadhawan, H. Bansal, K.-W. Chang, and N. Peng, "ConTextual: Evaluating context-sensitive text-rich visual reasoning in large multimodal models," 2024, *arXiv:2401.13311*.
- [234] H. Sun et al., "AdaPlanner: Adaptive planning from feedback with language models," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 36, pp. 58202–58245, 2023.
- [235] J.-F. Yeh et al., "AED: Adaptable error detection for few-shot imitation policy," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 37, pp. 136805–136836, 2024.
- [236] G. Canal, C. Torras, and G. Alenyà, "Are preferences useful for better assistance: A physically assistive robotics user study," *ACM Trans. Hum.-Robot Interact.*, vol. 10, no. 4, pp. 1–19, Dec. 2021.
- [237] R. Kachouie et al., "Socially assistive robots in elderly care: A mixed-method systematic literature review," *Int. J. Hum.-Comput. Interact.*, vol. 30, no. 5, pp. 369–393, May 2014.
- [238] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, "Bc-z: Zero-shot task generalization with robotic imitation learning," in *Conference on Robot Learning*, 2022, pp. 991–1002.
- [239] Y. Zhang et al., "Vision-and-language navigation today and tomorrow: A survey in the era of foundation models," 2024, *arXiv:2407.07035*.
- [240] F. Flammini et al., "Towards trustworthy autonomous systems: Taxonomies and future perspectives," *IEEE Trans. Emerg. Topics Comput.*, vol. 12, no. 2, pp. 601–614, Apr. 2024.
- [241] L. Chen et al., "Human-in-the-loop consensus tracking control for UAV systems via an improved prescribed performance approach," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 59, no. 6, pp. 8380–8391, Dec. 2023.
- [242] G. Lin, H. Li, C. K. Ahn, and D. Yao, "Event-based finite-time neural control for human-in-the-loop UAV attitude systems," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 34, no. 12, pp. 10387–10397, Dec. 2023.
- [243] J. Pesonen et al., "Detecting wildfires on UAVs with real-time segmentation trained by larger teacher models," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, Feb. 2025, pp. 5166–5176.
- [244] S. Javaid, H. Fahim, B. He, and N. Saeed, "Large language models for UAVs: Current state and pathways to the future," *IEEE Open J. Veh. Technol.*, vol. 5, pp. 1166–1192, 2024.
- [245] X. Zhou et al., "Swarm of micro flying robots in the wild," *Sci. Robot.*, vol. 7, no. 66, May 2022, Art. no. eabm5954.
- [246] H. Qin et al., "JD.Com: Operations research algorithms drive intelligent warehouse robots to work," *INFORMS J. Appl. Analytics*, vol. 52, no. 1, pp. 42–55, Jan. 2022.
- [247] R. Bogue, "The role of robots in logistics," *Ind. Robot: Int. J. Robot. Res. Appl.*, vol. 51, no. 3, pp. 381–386, May 2024.
- [248] I. F. A. Prawira et al., "Robot as staff: Robot for Alibaba e-commerce warehouse process," in *Proc. Int. Conf. Inventive Comput. Technol. (ICICT)*, Apr. 2023, pp. 1619–1623.
- [249] T. Jin and X. Han, "Robotic arms in precision agriculture: A comprehensive review of the technologies, applications, challenges, and future prospects," *Comput. Electron. Agricult.*, vol. 221, Jun. 2024, Art. no. 108938.
- [250] H. Xue et al., "Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation," 2025, *arXiv:2503.02881*.
- [251] T. Belpaeme et al., "Social robots for education: A review," *Sci. Robot.*, vol. 3, no. 21, Aug. 2018, Art. no. eaat5954.
- [252] G. Lampropoulos, "Social robots in education: Current trends and future perspectives," *Information*, vol. 16, no. 1, p. 29, Jan. 2025.
- [253] G. Wang et al., "Voyager: An open-ended embodied agent with large language models," 2023, *arXiv:2305.16291*.
- [254] L. Qunzhi et al., "Research on the application of embodied intelligence technology in space exploration," in *Proc. IEEE Int. Conf. Control Sci. Syst. Eng. (ICCSSE)*, Oct. 2024, pp. 149–155.