

Transforming Monolithic Foundation Models into Embodied Multi-Agent Architectures for Human-Robot Collaboration

Nan Sun¹, Bo Mao², Yongchang Li¹, Chenxu Wang¹, Di Guo² and Huaping Liu^{1,†}

Abstract—Foundation models have become central to unifying perception and planning in robotics, yet real-world deployment exposes a mismatch between their monolithic assumption that a single model can handle all cognitive functions and the distributed, dynamic nature of practical service workflows. Vision-language models offer strong semantic understanding but lack embodiment-aware action capabilities while relying on hand-crafted skills. Vision-Language-Action policies enable reactive manipulation but remain brittle across embodiments, weak in geometric grounding, and devoid of proactive collaboration mechanisms. These limitations indicate that scaling a single model alone cannot deliver reliable autonomy for service robots operating in human-populated settings. To address this gap, we present InteractGen, an LLM-powered multi-agent framework that decomposes robot intelligence into specialized agents for continuous perception, dependency-aware planning, decision and verification, failure reflection, and dynamic human delegation, treating foundation models as regulated components within a closed-loop collective. Deployed on a heterogeneous robot team and evaluated in a three-month open-use study, InteractGen improves task success, adaptability, and human-robot collaboration, providing evidence that multi-agent orchestration offers a more feasible path toward socially grounded service autonomy than further scaling standalone models.

Index Terms—Foundation models, multi-agent systems, service robotics, human-robot collaboration.

I. INTRODUCTION

FOUNDATION models have rapidly entered the robotics community as a unifying abstraction for perception, language, and action [1], [2], [3], [4], yet their real-world impact on service robots remains far from the autonomy required in human-populated environments. Current robotics foundation models fall into two powerful but fundamentally incomplete families. The first comprises *embodied Vision-Language Models (VLM)*—such as VeBrain [5], RoboBrain [6] and MiMo-Embodied [7]—which unify multimodal understanding and spatial reasoning, and often demonstrate strong performance on planning. However, these models typically depend on post-training with benchmark-specific adaptations, in which abilities distilled from larger VLMs are operationalized by generating training data using diverse task templates; however, their action

components are generally delegated to hand-designed low-level skill libraries, as in SayCan [8], and concrete demonstrations of multi-robot execution, embodiment-aware action, or human-centered deployment remain limited.

The second family comprises *Vision-Language-Action (VLA) policies*, which primarily fine-tune relatively smaller VLMs on robot action datasets using autoregressive language heads [9], [10] or diffusion/flow-based action experts [11], [12], [13], achieving promising reactive manipulation across multiple embodiments and scenes. However, deploying a model in *real-world service settings* requires *comprehensive autonomy*: an assistant operating in such environments must simultaneously parse instructions, generate appropriate motion plans, manage uncertainty, and interact fluidly with surrounding human users. Under these demands, current VLAs serve as excellent motor primitives but remain inadequate as assistive systems: they lack persistent task memory, contextual reflection, collaboration awareness, and mechanisms for interacting with humans within realistic workflows. Together, these trends highlight a structural tension: despite substantial isolated capabilities, neither family of foundation models yet realizes the intentional, contextually grounded autonomy required for service robots to deliver meaningful social utility.

These considerations point toward a broader insight: the difficulty of deploying current foundation models in real-world service robots does not arise merely from limited scale or training data, but from a deeper mismatch between their monolithic architectures and the distributed, interdependent nature of actual service tasks. Real-world environments inherently involve multiple humans, robots, tools, and information channels, and therefore require cognitive structures capable of arbitrating uncertainty and coordinating heterogeneous embodiments over long horizons. Evidence from digital multi-agent LLM systems such as AutoGen [14], MetaGPT [15], and AgentVerse [16] reinforces this view: even without embodiment, reliable long-horizon autonomy typically emerges from role decomposition, reflective loops, tool invocation, and explicit verification steps, rather than from a single large model. When physical robots enter the loop—with real-world dynamics, partial observability, safety constraints, and human unpredictability—the limitations of monolithic foundation models become even more pronounced. For the foreseeable future, it is unlikely that any single VLM or VLA claiming to be an embodied foundation model will simultaneously deliver spatial grounding, cross-

¹The author is with the Department of Computer Science and Technology, Tsinghua University, Beijing, 100084, China.

²The author is with the School of Artificial Intelligence, Beijing University of Posts and Telecommunications, Beijing, 100876, China.

[†]Corresponding Authors. hpliu@tsinghua.edu.cn




Fig. 1. A holistic demonstration of how InteractGen coordinates robots and humans in real time. The framework monitors online and offline environments, performs collaborative reasoning across specialized agents, plans composite workflows, and triggers physical execution through heterogeneous robots. Humans are treated as deployable agents—InteractGen clarifies, notifies, and delegates subtasks when appropriate—enabling socially grounded service autonomy.

embodiment skill composition, multi-agent task allocation, safety arbitration, and socially appropriate interaction. What is needed instead is an architectural shift—one that treats foundation models not as monolithic controllers but as powerful *components* embedded within structured multi-agent control systems that support *collaborative embodied reasoning*. Such a design enables the flexibility, interpretability, and competence required for a genuinely general-purpose robotic brain.

Motivated by this perspective, we introduce InteractGen, an LLM-powered multi-agent framework enabling collaborative embodied reasoning through specialized agents built on diverse foundation models (Fig. 1). InteractGen coordinates heterogeneous robots and humans via five agents for perception, planning, decision making, validation, and reflection. This decomposition extends beyond single-loop paradigms such as ReAct [17] and Reflexion [18], while avoiding decentralized designs requiring a separate LLM per robot as in EMOS [19] which is hard to scale. InteractGen can also invoke API-level tool calls for GUI operations while serving as a unified “robot brain” for low-level skill libraries or end-to-end VLA policies. By centralizing high-level reasoning and delegating physical execution to generalist controllers or modular primitives, InteractGen treats foundation models as collaborative components rather than monolithic controllers.

Crucially, InteractGen is designed as a human-centered interactive system that takes humans as *deployable agents* capable of taking on delegated subtasks whenever robots encounter safety, authority, or dexterity limits. This allows the system to recognize its own bounds and proactively solicit human assistance when tasks exceed robotic capability in dynamic environments, pushing embodied applications beyond isolated navigation or manipulation toward realistic service workflows that require joint digital–physical coordination. Across controlled evaluations and a three-month open-use de-

ployment, we observed that even simple mobile robots became markedly more useful when embedded in this human-aware loop—acting not merely as executors but as assistants that could anticipate needs, delegate tasks when appropriate, and engage in extended multi-user workflows. These results suggest that multi-agent orchestration, combined with meaningful human involvement, offers a more promising path to socially grounded autonomy than further scaling monolithic foundation models, underscoring that effective human–robot collaboration is central to the future of service robotics.

In summary, this work makes four key contributions:

- **Insights into designing multi-agent frameworks built on foundation models.** We identify structural weaknesses in single-model robotic foundations and demonstrate how different classes of foundation models can be systematically embedded into multi-agent architectures tailored for embodied tasks, providing design principles for scalable, interpretable, and robust embodied systems.
- **An LLM-driven multi-agent architecture for real-world multi-user tasks.** We introduce **InteractGen**, a framework that decomposes robot intelligence into specialized agents for perception, planning, decision making, verification, and reflection—treating foundation models as modular components rather than monolithic controllers to enable collaborative embodied reasoning.
- **A novel human-as-deployable-agent paradigm for effective Human-Robot Collaboration.** We formalize a paradigm in which humans serve as *deployable agents* who can be delegated subtasks when robots face safety, authority, or dexterity limits, enabling robust teaming in realistic multi-user environments with dynamic human participation and hybrid digital–physical workflows.
- **Large-scale real-world deployment demonstrating generalization and social utility.** We deploy InteractGen

TABLE 1

COMPARISON BETWEEN INTERACTGEN AND PREVAILING AGENTIC SYSTEMS. *Integrated* DENOTES SINGLE-STEP ACTION GENERATION, WHEREAS *Hierarchical* DENOTES EXPLICIT SEPARATION BETWEEN PLANNING AND DECISION MAKING. *Validation* REFERS TO PRE-EXECUTION SELF-CHECKING, AND *Reflection* REFERS TO POST-EXECUTION REASONING. IN THE **INTERACTION SKILLS** COLUMNS, GREY RECTANGLES INDICATE THAT THE SYSTEM DOES NOT EXPLICITLY SUPPORT THIS CAPABILITY IN ITS DESIGN OR OPERATION.

Methods	REASONING COMPOSITION				INTERACTION SKILLS		
	Perception	Planning & Decision	Validation	Reflection	Robot Dispatch	GUI Operation	Human Collaboration
ReAct [17]	✓	<i>Integrated</i>	✗	✗	●	□	□
Reflexion [18]	✓	<i>Integrated</i>	✗	✓	●	□	□
Mobile-Agent-v2 [20]	✓	<i>Hierarchical</i>	✗	✓	□	●	□
KnowNo [21]	✓	<i>Integrated</i>	✓	✗	●	□	<i>Clarification</i>
SMART-LLM [22]	✓	<i>Integrated</i>	✗	✗	●	□	□
CoELA [23]	✓	<i>Hierarchical</i>	✓	✗	●	□	□
CaPo [24]	✓	<i>Hierarchical</i>	✓	✗	●	□	□
RoCo [25]	✓	<i>Integrated</i>	✗	✗	●	□	<i>Supervision</i>
Lip-LLM [26]	✓	<i>Hierarchical</i>	✓	✗	●	□	□
EMOS [19]	✓	<i>Integrated</i>	✓	✗	●	□	□
LaMMA-P [27]	✓	<i>Hierarchical</i>	✓	✗	●	□	□
HMCF [28]	✓	<i>Hierarchical</i>	✓	✓	●	□	<i>Supervision</i>
InteractGen (Ours)	✓	<i>Hierarchical</i>	✓	✓	●	●	<i>Subtask Delegation</i>

across a heterogeneous robot team and conduct extensive experiments plus a three-month open-use study with non-expert occupants, demonstrating strong cross-embodiment generalization, reliable human-robot collaboration, and substantial social value in real-world service environments.

II. RELATED WORK

A. Foundation Models for Robotics

Foundation Models for Robotics are typically seen as large multimodal backbones supporting perception, language understanding, and action. Existing approaches largely follow two trajectories. VLMs such as CLIP [29] and PaLI-Gemma [30] provide strong scene understanding and semantic grounding. Embodied variants including VeBrain [5], RoboBrain [6], and MiMo-Embodied [7] extend these capabilities into physical environments, often through distillation and skill-conditioned prompting, as in SayCan [8]. That being said, existing VLMs largely rely on hand-crafted controllers and exhibit limited robustness in action grounding.

VLA policies such as RT-2 [9], π_0 [11], OpenVLA [10], and CogAct [12] augment VLMs with end-to-end action heads. Despite improved perception–action coupling, these policies remain reactive, with limited task memory, weak uncertainty modeling, and no mechanisms for human involvement or multi-agent coordination. Even in structured VLN and VLA benchmarks, both VLMs and VLAs remain fundamentally single-agent systems.

In summary, current foundation models fall short of the distributed, interactive cognition required in multi-user, real-world service environments, motivating their integration as modular components within more comprehensive frameworks.

B. LLM-Powered Multi-Agent Frameworks

Recent work in virtual domains increasingly adopts LLM multi-agent architectures, where specialized agents coordinate

via explicit role decomposition. Frameworks such as AutoGen [14], MetaGPT [15], and AgentVerse [16] show advantages in long-horizon reasoning and verification over single-agent methods such as ReAct [17] and Reflexion [18]. Yet these systems remain fundamentally disembodied, without grounding in perception, physical constraints, or actuation.

While such virtual-domain systems highlight the importance of structured role decomposition, their embodied counterparts remain far more [31], [32]. Existing attempts toward embodied multi-agent systems make only partial progress. SMART-LLM [22] integrates multiple modules but lacks explicit role structure; CoELA [23] and CaPo [24] support multi-robot collaboration but overlook human interaction and hierarchical verification; LaMMA-P [27] focuses on simulated multi-agent planning without real-world actuation.

In contrast, InteractGen unifies specialized foundation models across five reasoning roles and couples them with physical execution pipelines and human-facing interfaces. This enables embodied multi-agent coordination that dispatches robots, uses GUI tools, verifies execution, and dynamically delegates subtasks to humans, offering broader reasoning coverage and richer interaction skills than prior work (see Table 1).

C. Human–Robot Collaboration in Agentic Systems

Human–robot collaboration in agentic systems typically follows two paradigms. The human-as-clarifier model involves humans resolving ambiguities during execution [21]. Although effective for disambiguation, this restricts humans to simple query answering and becomes inefficient in long-horizon embodied tasks [33]. The human-as-supervisor paradigm places humans in high-level oversight roles, but continuous monitoring imposes substantial cognitive load [25].

A third emerging paradigm treats humans as deployable agents for subtasks beyond robotic capability. Early systems such as CoBots [34] and modern assistants like AssistantX [35] demonstrate proactive human involvement, but existing im-




Fig. 2. **Overview of the InteractGen architecture.** InteractGen naturally supports three operating modes that emerge to handle long-horizon, human-centered scenarios. **REACTIVE:** Perceiver perceives task-relevant information and signals Planner to generate a Thought-of-Action plan; Assigner assigns actions to suitable agents. **ACTIVE:** Manager triggers clarification for ambiguous cases and Validator’s validation elicits re-planning, avoiding rigid generate-then-execute patterns. **PROACTIVE:** Manager reflects, corrects prior reasoning, and reactivates Perceiver for proactive reasoning. These modes enable interactive coordination with humans and robots in dynamic environments. The *act-fail-reflect-replan* mechanism in Proactive Mode greatly enhances human-robot collaboration.

plementations remain narrow, rule-based, or domain-specific. InteractGen advances this paradigm by explicitly integrating humans into the multi-agent loop as active collaborators: robots can delegate subtasks and request human actions for tasks they cannot safely, legally, or dexterously perform, thereby achieving better overall task outcomes. This enables more natural, scalable, and socially grounded teamwork in dynamic, multi-user service workflows where robot-only autonomy is insufficient.

III. METHODOLOGY OF INTERACTGEN

A. Agent Specifications of InteractGen

InteractGen consists of five LLM-driven agents operating over a shared message pool and a persistent world-state memory (see Fig. 2). Details of the choices on foundation models for each agent, together with the corresponding insights, are further presented in Section VI and Section VIII.

1) *Agent Manager (A_{mgr}) – Reflection & Clarification:* A_{mgr} is the reflection agent, whose output $\mathcal{RF}_t$ consists of three components: (i) a binary success flag Y/N indicating whether the previous plan $\mathcal{PL}_{t-1}$ and execution σ_{t-1} produced the desired outcome given perception $\mathcal{PC}_t$; (ii) a reflection thought explaining the result and providing corrective guidance; and (iii) a concise summary of $\mathcal{M}_{t-1}$. If the flag is Y , A_{mgr} signals A_p to proceed with planning; if N , it activates A_c to gather additional context $\Delta\mathcal{PC}_t$, enabling proactive re-planning. As illustrated in Fig. 3, the Manager’s self-reflection and flexible communication allow the system to adapt to real-world disturbances and autonomously revise the plan for robust execution. Its process can be formalized as:

$$(\mathcal{I}, \mathcal{PL}_{t-1}, \sigma_{t-1}, \mathcal{M}_{t-1}, \mathcal{PC}_t) \xrightarrow{A_{mgr}} \mathcal{RF}_t \quad (1)$$

A_{mgr} is also equipped with a *clarification module*, invoked upon receiving the initial user instruction, that returns either {clear} or {unclear + clarification question}; if the result is unclear, the system relays the

question directly to the user for active clarification, as shown in Fig. 4.

2) *Agent Perceiver (A_c) – World-State Perception:* A_c receives prior world state $\mathcal{W}_{t-1}$, plan $\mathcal{PL}_{t-1}$, and executed actions σ_{t-1} (including completion signals from robots), and produces perception output $\mathcal{PC}_t$ while updating $\mathcal{W}_{t-1}$ to $\mathcal{W}_t$. $\mathcal{PC}_t$ encodes structured representations of task-relevant entities and environmental conditions for iteration t . If A_{mgr} reports failure, A_c is prompted to invoke perceptrors from the integrated *PerceptorHub*—which unifies modalities such as LiDAR-based localization, RealSense vision, and online dialogue or GUI APIs—to produce reflection-driven, real-time proactive observations $\Delta\mathcal{PC}_t$ for re-planning. The operation of this agent can be formalized as:

$$(\mathcal{I}, \mathcal{W}_{t-1}, \mathcal{PL}_{t-1}, \sigma_{t-1}) \xrightarrow{A_p} \mathcal{PC}_t \quad (2)$$

3) *Agent Planner (A_p) – Thought-of-Action Planning:* Given the current perception $\mathcal{PC}_t$ and reflection $\mathcal{RF}_t$, the planner A_p produces a *Thought-of-Action* plan $\mathcal{PL}_t$: a reasoning trace paired with an ordered sequence of sub-tasks required to accomplish $\mathcal{I}$. By hierarchically decoupling planning from execution, A_p —trained through a three-stage pipeline to emulate human problem-solving via skill decomposition—identifies executable, dependency-free actions and arranges them into a coherent, interpretable roadmap, such as: 1. `MoveTo(printer)`, 2. `Print(e-file)`, 3. `DeliverTo(Alice)`. The corresponding transformation is:

$$A_p: (\mathcal{I}, \mathcal{PC}_t, \mathcal{RF}_t) \xrightarrow{A_p} \mathcal{PL}_t \quad (3)$$

4) *Agent Assigner (A_s) – Robot-Human Coordination:* A_s consumes the ToA plan from A_p , the robot profiles and skill library, and active contacts/groups from $\mathcal{PC}_t$, alongside few-shot prompts covering both robot-only and robot-human hybrid scenarios. By eliciting latent language understanding from the LLM, A_s produces clear assignments (e.g., `assign(MoveTo(Printer), Robot1)` or




Fig. 3. A proactive case of InteractGen reasoning and coordination. The system executes the instruction through a modular pipeline. The PerceptorHub collects raw sensor and chat signals, which are incrementally updated by the Perceiver. The Planner generates dependency-aware Thought-of-Action steps, while the Manager monitors progress and triggers reflection when inconsistencies arise (e.g., Carl not at his usual location). The Assigner distributes subtasks to robots or humans, and the Validator checks feasibility before execution. This act-reflect-replan loop enables InteractGen to adapt to dynamic human availability and environmental changes, ensuring reliable multi-user task execution in real-world settings.

`request(Alice, Printing(loc=Printer))`, maximizing autonomous execution and deferring to humans only when robots are incapable. The agent's input-output mapping is defined as follows:

$$A_s : (\mathcal{I}, \mathcal{PC}_t, \mathcal{RF}_t, \mathcal{PL}_t, \mathcal{W}_t) \xrightarrow{A_s} \sigma_t. \quad (4)$$

5) **Agent Validator (A_v) – Pre-Execution Verification:** A_v uses a dedicated LLM as a judge to verify σ_t against $\mathcal{I}$, $\mathcal{W}_t$, and the summary of prior execution in $\mathcal{RF}_t$, rather than relying on a continuation of the planner's or assigner's chain-of-thought to rethink, thereby reducing bias accumulation. $\mathcal{V}_t$ contains (i) a validity flag `pass/fail` assessing the correctness of σ_t , and (ii) a brief diagnostic thought explaining the verdict. If invalid, the explanation guides A_p 's replanning by identifying specific mismatches. This validation enhances robustness in human-interactive tasks where execution errors are costly. Formally, the agent computes the following transformation:

$$A_v : (\mathcal{I}, \mathcal{W}_t, \mathcal{RF}_t, \sigma_t) \xrightarrow{A_v} \mathcal{V}_t \quad (5)$$

B. Memory Unit

The Memory Unit serves as the foundational layer of the entire framework. It maintains both *long-term* and *short-term* memories (see Fig. 5). Short-term memory stores the transient information generated during each end-to-end execution of a user request, including the instruction, recent dialogue context, and embodied states such as the robot's location or the smart locker's status. The agent's textual outputs retained in the shared message pool also function as part of short-term memory, enabling rapid within-session reasoning without repeatedly invoking perception modules.

Long-term memory, in contrast, is modeled as an undirected topological graph. Nodes correspond to humans, public facilities, personal items, and physical locations, while edges encode their relationships—for example, item ownership or the spatial placement of humans and facilities. Human nodes additionally maintain an availability attribute indicating their readiness to participate in collaborative tasks. During execution, the Perceiver incrementally updates node attributes and relations




Fig. 4. Example of active reasoning in InteractGen. The Manager raises a clarification question when input is ambiguous. After validation fails, the Planner re-generates a feasible plan using updated context. This showcases the system’s ability to actively adapt before execution errors occur.




Fig. 5. Overview of the Memory Unit. The left panel illustrates the long-term memory graph, which incrementally encodes cross-task knowledge, entity relations, and reusable dependencies accumulated throughout past interactions. The right panel shows the short-term memory, which stores task-specific observations, intermediate reasoning states, and execution feedback to support real-time planning and adaptation.

to reflect environmental changes, ensuring that the memory remains concise and semantically grounded rather than a raw, ever-growing log of observations.

We denote the overall memory state at time step t as $\mathcal{M}_t$, which can be serialized into textual descriptions and supplied to the foundation models as part of their contextual input. Combined with the Perceiver’s refinement of memory entries, this mechanism reduces the token footprint, improves accuracy by filtering out irrelevant noise, and enables foundation models to perform faster and more stable reasoning for long-horizon interactive tasks. By maintaining only on-demand, structured information, and by combining the Perceiver with this memory design, the system greatly lightens the contextual burden passed to downstream agents: the Manager, Planner, Assigner, and Validator operate on compact, semantically meaningful summaries rather than large, unfiltered histories.

IV. DATASET CONSTRUCTION

Constructing reliable embodied reasoning agents requires evaluation settings that reflect the complexity of real-world service robotics, where multiple humans, robots, and shared resources must be coordinated under evolving constraints. Existing benchmarks largely isolate agents from human participants or restrict interaction to static manipulation, leaving a substantial gap to practical deployment. To bridge this gap, we build a high-quality dataset integrating natural-language understanding, instruction clarification, and parallel task allocation within human-centric workflows. The dataset contains

2,100 task instances from simulated domestic and workplace environments and is organized into three tiers: (1) *Basic flows* (300) with fully specified requests; (2) *Ambiguous flows* (800) requiring clarification of roles, targets, or intents; (3) *Dynamic-context flows* (1,000) involving human unavailability or resource contention that demand adaptive planning.

1) *Basic Instruction Generation*: Each task is generated from a structured JSON schema describing multi-step interactions among humans, public facilities, and personal items. Personal items are explicitly grounded to owners, while public facilities (e.g., printers, lockers) encode functional constraints. We use DeepSeek-R1 to sample grounded entities from the environment, assigns semantic roles—origin, intermediary, and destination users—and generates dependency-aware action sequences forming executable service workflows, with irrelevant entities pruned automatically.

All tasks share a controlled vocabulary of 30 portable personal items and 12 public facilities, covering representative object types and functional resources commonly required in human-robot collaborative scenarios.

Personal Item	Mass (g)	Bbox (cm)
Ballpoint_Pen	10	14×1×1
Stapler	220	13×4×5
Document	5	29.7×21×0.05
...		
Public Facility	Purpose / Interaction	
Printer	Print, scan, copy documents	
Fridge	Store perishable items	
Coffee_Machine	Support SelectDrink	
...		

The resulting JSON structure reflects both the role-specific object associations and the dependency-aware action graph:

```
{
  "origin": ["person_who_issued_request"],
  "destination": ["final_recipient"],

  "facility": ["required_facility"],
  "facility_usage_action": ["action_to_use_facility"],

  "personal_item_1": ["item_name"],
  "personal_item_1_owner": ["item_owner"],
  "personal_item_1_actions": [
    ["action", "person", "dependency(optional)"],
    ...
  ],
  "intermediate_person_1": [
    "person_involved_in_transit_or_signature"],
    ...
}
```

This structure ensures that every entity is role-assigned and every action is grounded to an actor and, when necessary, a prerequisite. Natural-language instructions are generated from the structured JSON and validated through a semi-automated pipeline: 83% of samples pass automatic checks, and the remaining 17% are manually corrected, with duplicates removed. The resulting set contains 300 high-quality, dependency-aware task specifications serving as the basic service flows.

2) *Ambiguity Injection*: To simulate underspecified real-world requests, we convert each basic task into an ambiguous variant by masking one critical semantic field (e.g., facility

usage, personal item, required actor, or destination). Few-shot prompting with GPT-4o then produces an ambiguous instruction together with a corresponding clarification question, generated from both the complete and masked JSON. This yields 800 ambiguity-driven cases.

3) *Dynamic Context Simulation*: To model realistic disruptions, we augment each task with a `dynamic_context` field specifying time-varying human availability. A subset of replaceable actors is randomly marked as unavailable; this information is hidden from the planner and used only by the simulator to trigger failures when unavailable individuals are referenced. These failures require the planner to revise its strategy, enabling evaluation under partial observability, imperfect inputs, and runtime recovery. This stage adds 1,000 dynamic-context tasks.

V. THOUGHT-OF-ACTION PLANNING

We introduce a *Thought-of-Action* (ToA) representation that bridges free-form chain-of-thought reasoning and executable primitive skills. Each action is typed, parameterized, and dependency-aware, enabling parallelism and precondition checking. Compared with generic chain-of-thought, ToA provides a structured, dependency-aware representation that enables the Assigner to dispatch actions to humans or robots and allows the Validator to check feasibility before execution. This compact formulation significantly reduces token usage and improves planning efficiency.

Our planner is trained through a three-stage pipeline (Fig. 6):

1) *Stage 1: Imitation Learning of Reasoning*: We fine-tune an LLM (Qwen3-8B in InteractGen) on 5,000 supervised examples constructed from 1,500 tasks randomly sampled from Section IV. Each sample is augmented with intermediate perception and reflection states, allowing the planner to learn mid-execution reasoning rather than cold-start decomposition. Manually annotated, dependency-aware plans, which can be seen as a form of simple chain-of-thought, serve as supervision targets.

Given $\mathcal{D} = \{(x_i, y_i)\}$ where x_i concatenates all context fields and y_i is the plan token sequence, the IL objective is:

$$\mathcal{L}_{IL} = -\frac{1}{N} \sum_{i,t} \log p_{\theta}(y_{i,t} \mid y_{i,<t}, x_i). \quad (6)$$

This establishes strong priors for coherent, dependency-aware reasoning.

2) *Stage 2 – Thought-to-Action Grounding*: Building on STAGE 1, we further refine the model by aligning its free-form reasoning with executable skill primitives. We introduce a human-atomic skill library across public domains, including actions like `walkto(location)`, `pickup(object)`, `print(document)`, etc. Each training sample maps an instruction to a structured *Thought-of-Action* plan—a sequence of minimal skills from the library. For example, “Alice: Print a report in the USB Drive for the meeting in 502 hosted by Bob” is grounded as: `WalkTo(Alice)`, `Retrieve(USB Drive)`, `WalkTo(Printer)`, `Insert(USB Drive)`, `Print(Report)`, `DeliverTo(Bob, Room 502)`.




Fig. 6. **The three-stage training pipeline of ToA Planning.** The pipeline progressively transforms natural-language instructions into structured, executable Thought-of-Action plans. **Stage 1: Imitation Learning of Reasoning.** The model learns to imitate human planning demonstrations and acquire global dependencies between reasoning steps. **Stage 2: Thought-to-Action Grounding.** Using GRPO rollouts, the model is trained to abstract fine-grained action skills, with rewards encouraging correct action formats, parameters, and ordering. **Stage 3: Rejection Sampling and Refinement.** The planner is further refined via SFT and rejection sampling, where LLM-based judges filter incorrect or low-quality outputs to enhance correctness, robustness, and generalization.

a) *Human-Atomic Skill Library:* We focus on *human-atomic* skills that robots may not yet master fully but that humans can readily perform. Each skill record includes following components:

```
SkillName(parameters), e.g. MoveTo(Alice),
Order: (step = n, dep = k), where dep denotes the
most recent prerequisite step.
```

b) *Dataset Augmentation for GRPO:* We create 5,000 *skill-level trajectories* by pairing every STAGE 1 context with its aligned skill sequence. Two strategies were evaluated:

- 1) *Expert Annotation:* human annotators map the gold plan to skills—high accuracy but costly.
- 2) *Model Bootstrapping:* the STAGE 1 planner auto-generates a plan, which is then heuristically parsed into skills and provides on-policy diversity.

Empirically we combine both: 60% expert-labelled seeds ensure correctness; 40% model-generated samples broaden coverage and expose current stage to its own prediction distribution. The data format is as follows:

```
{
  "instruction": "user instruction",
  "input": {
    "env info": "prior knowledge",
    "perception": "perception of last action",
    "reflection": ["Y/N", "description", "history"],
  },
  "output": {
    "planning with dependency": "planning"
    "action with dependency": "human-atomic skills"
  }
}
```

c) *Reward Design for GRPO:* To optimize the planner, we adopt GRPO—a group-based variant of PPO that evaluates multiple sampled action sequences per query. Since GRPO requires a scalar reward r_i for each predicted sequence s_i , we design a reward function that captures both structural validity

and semantic correctness, measuring how closely each predicted skill sequence aligns with the ground-truth execution plan.

Formally, the reward evaluates four aspects:

Format validity. We first check whether the output is structurally valid—i.e., whether the required fields (planning dependencies, skill sequence, parameters) appear in the correct format:

$$\mathcal{R}_{\text{format}} = \begin{cases} 1, & \text{all fields are present and formatted} \\ 0, & \text{otherwise} \end{cases} \quad (7)$$

Skill-name matching. Let N_G and N_P be the sets of skill names appearing in the ground-truth sequence G and predicted sequence P . Their similarity is measured using the Jaccard index:

$$r_{\text{name}} = \frac{|N_G \cap N_P|}{|N_G \cup N_P|}, \quad r_{\text{name}} \in [0, 1] \quad (8)$$

Parameter matching. For each ground-truth skill G_j , we compute the Jaccard similarity of its parameter set against the matched predicted skill P_j :

$$r_{\text{param}} = \frac{1}{|G|} \sum_{G_j \in G} \frac{|\text{params}(G_j \cap P_j)|}{|\text{params}(G_j \cup P_j)|}, \quad r_{\text{param}} \in [0, 1] \quad (9)$$

Ordering consistency. We further evaluate whether the predicted sequence preserves necessary action dependencies and subtask ordering:

$$r_{\text{order}} = \frac{1}{\sum_{G_j \in G} |\text{orders}(G_j)|} \sum_{G_j \in G} \sum_{o \in o(G_j)} \frac{1}{2} \mathbf{1}[P_j[o] = G_j[o]], \quad r_{\text{order}} \in [0, 1] \quad (10)$$

Combined matching score. The three semantic scores are summed to obtain an overall matching value:

$$r_{\text{match}} = r_{\text{name}} + r_{\text{param}} + r_{\text{order}}, \quad r_{\text{match}} \in [0, 3] \quad (11)$$

We then linearly normalize this into the correctness reward:

$$\mathcal{R}_{\text{correct}} = 2 \cdot r_{\text{match}} - 3, \quad \mathcal{R}_{\text{correct}} \in [-3, 3] \quad (12)$$

Final reward. The final reward used by GRPO is:

$$\mathcal{R}_{\text{final}} = \mathcal{R}_{\text{format}} + \mathcal{R}_{\text{correct}}, \quad \mathcal{R}_{\text{final}} \in [-3, 4] \quad (13)$$

For each query Q , GRPO samples a group $\mathcal{G}_Q = \{A, (s_1, r_1), \dots, (s_n, r_n)\}$, where A is the ground-truth plan and each s_i is a predicted sequence assigned reward r_i .

We compute the group mean and standard deviation:

$$\mu_Q = \frac{1}{n} \sum_{i=1}^n r_i, \quad \sigma_Q = \sqrt{\frac{1}{n} \sum_{i=1}^n (r_i - \mu_Q)^2} \quad (14)$$

The normalized advantage of each sample is:

$$A_i(s_i | Q) = \frac{r_i - \mu_Q}{\sigma_Q + \eta} \quad (15)$$

where η is a small constant for numerical stability. The final GRPO optimization objective can be represented as:

$$J_{\text{GRPO}}(\theta) = \mathbb{E}_Q \left[\min \left(\frac{\pi_\theta(s_i | Q)}{\pi_{\text{old}}(s_i | Q)} A_i, \right. \right. \\ \left. \left. \text{clip} \left(\frac{\pi_\theta(s_i | Q)}{\pi_{\text{old}}(s_i | Q)}, 1 - \epsilon, 1 + \epsilon \right) A_i \right) \right] \quad (16)$$

where π_θ and π_{old} denote the current and reference policies, and ϵ controls the clipping range to prevent overly large policy updates.

Our group-normalized objective stabilizes training by reducing reward variance across queries and provides a much finer supervision signal than binary or coarse scoring. By evaluating skill names, parameters, and ordering jointly, the reward reflects the true structure of tool-using tasks, leading to plans that are more consistent and executable. In practice, this yields faster convergence and stronger performance than imitation learning alone. GRPO-trained planners produce cleaner, dependency-correct action sequences and are less prone to formatting or semantic errors, while also learning to request clarification when uncertain.

3) *Stage 3: Rejection Sampling and Refinement:* We expand the task pool with 500 additional instructions (300 curated + 200 adapted from ALFWorld [36] and BEHAVIOR [37]). From the GRPO-trained planner, we sample $K = 10$ diverse ToA candidates per task. A high-capacity verifier (Gemini 2.5 Pro) filters candidates based on logical consistency, semantic validity, and conciseness. A total of 8,123 trajectories are accepted.

Each accepted trajectory $\tau = (y, a)$ supervises both planning and action decoding:

$$\mathcal{L}_{\text{RST}}(\theta) = -\frac{1}{|\mathcal{T}_{\text{acc}}|} \sum_{(y,a) \in \mathcal{T}_{\text{acc}}} \left[\sum_{t=1}^{|y|} \log p_\theta(y_t | y_{<t}, x) \right. \\ \left. + \sum_{t=1}^{|a|} \log p_\theta(a_t | a_{<t}, y, x) \right] \quad (17)$$

This final refinement distills verifier-approved reasoning patterns, yielding a planner that is accurate, dependency-consistent, and robust in real-world workflows.

VI. EXPERIMENT

Our evaluation targets service-oriented workflows that couple item operations, user localization, and human–robot interaction subtasks (e.g., requesting signatures, borrowing objects, performing handovers) in fluid, multi-user settings. Experiments are conducted in two parallel environments sharing identical robot profiles and coordination protocols. The simulation environment extends TextWorld [38] to support multi-user office scenes with symbolic robot skills and natural-language user responses. The real-world environment is a 400 m² lab–office (see Fig. 7) containing multiple users, 12 personal-item categories, and 6 shared facilities. Evaluation spans 300 unseen simulated tasks (five runs each) and 167 real-world tasks (three runs each, including ablations), with mirrored APIs and capabilities ensuring consistent comparison.

A. Robot Profiles

We deploy a heterogeneous team of four robot types, each contributing specialized mobility and manipulation capabilities. The Assigner allocates subtasks using structured robot profiles encoding affordances, workspace limits, sensing, and dexterity.

a) *Mobile Base with Smart Locker (two units):* A chassis-based mobile robot equipped with smart lockers handle autonomous delivery. An RGB-D camera and 2D LiDAR support navigation, while QR-code sensing provides reliable feedback on locker contents. Lacking external manipulation, they specialize in secure transport.

b) *Wheeled Robot with Rear Tray:* A lightweight shuttle platform carrying a passive rear tray. Basic cameras and navigation sensors support flexible intra-room transfers.

c) *Quadruped Robot with Manipulator:* A legged robot with an integrated arm for operation in cluttered, uneven, or low-elevation environments. Its manipulation controller is trained via the ACT framework [39] using 600 teleoperated demonstrations collected for tasks like pick, place, and handover. Data include synchronized RGB-D streams and end-effector poses obtained through a joystick–tablet teleoperation setup.

d) *Mobile Manipulator:* A wheeled base with a 6-DOF arm serving as the team’s most dexterous platform. It excels at mid- and high-elevation tasks (desk/shelf retrieval, drawer operations, handovers). Its ACT controller is trained on 800 teleoperated trajectories gathered via a VR interface, with synchronized joint states and RGB-D data.

e) *Human Agent:* Humans act as flexible collaborators for tasks requiring fine dexterity, domain knowledge, or capabilities beyond the robots. The Assigner issues symbolic requests (e.g., `ask_human(print, file)`), and humans respond through a dedicated messaging thread. InteractGen’s LLM agents run on a separate compute node and dispatch symbolic commands to robots through ROS. The Perception and Reflection agents maintain a concise, continuously updated semantic world state appended to each planning prompt.

B. Real-World Environment

Our real-world deployment takes place in a 400 m² lab–office environment designed to emulate realistic domestic scenarios (see Fig. 7). The space is divided into functional zones




Fig. 7. **Overview of the experimental setup.** Our system operates across a heterogeneous fleet of embodied robots and an interactive real-world environment, accompanied by a mirrored simulation environment with a language interface. The top row shows the four robot platforms used in our study—mobile base, wheeled robot, quadruped, and mobile manipulator—highlighting their sensing and actuation capabilities. The bottom row illustrates the corresponding semantic maps, real-world scenes, and simulation interface that together constitute the operational environment for executing human–robot collaborative tasks.

(e.g., meeting room, workspace) and populated with 6 human participants, 12 personal-item categories, and 6 public facilities.

To support autonomous multi-robot planning and execution, we integrate *five types of perceptors* that provide multi-modal situational awareness. These perceptors interface with physical hardware (RGB-D cameras, LiDAR, force sensors) or cloud APIs (dialogue streams), continuously assisting with updating the unified memory unit shared across all agents.

Vision Perceptor (Detection and Scene Understanding). Each mobile robot carries an Intel RealSense D435 RGB-D camera. An RGB-based detector identifies household objects, while depth maps provide 3D approximations of their positions. When finer granularity is required, an image-segmentation model extracts object masks. Detected objects are anchored to the semantic map with spatial coordinates.

Localization and Navigation Perceptor. Robots localize using 2D LiDAR-based SLAM (RPLIDAR A3 with SLAMTEC SLAM), maintaining a shared (x, y, θ) frame across the heterogeneous fleet. A ROS Navigation Stack planner evaluates navigation feasibility and produces collision-free paths, allowing low-level mobility to be handled reliably and freeing our system to focus on interaction-centric reasoning for multi-user service workflows. The mobile robots combine LiDAR, wheel odometry, and an IMU to run a 2D pose-graph SLAM back-

end that generates a 5 cm occupancy grid. A semantic layer is built on top of this grid, annotating default locations of users, public facilities, and personal items through either automatic segmentation (e.g., Mask2Former) or manual annotation. The resulting metric–semantic map provides a compact set of task-relevant nodes that supports efficient reasoning about multi-user presence and navigation without requiring end-to-end vision–language navigation pipelines.

Force/Tactile Perceptor. Manipulator-equipped robots use a wrist-mounted force–torque sensor and gripper touch sensors to monitor grasping. Excessive forces or unexpected torque profiles signal grasp failure, while rapid finger-width reduction indicates successful pickup. Signals are processed at 50 Hz through ROS.

Dialogue Perceptor. A cloud-based interface monitors message threads and uses a retrieval API to stream new instructions, clarifications, and updates in real time, enabling language-driven multi-step planning. The system also registers as a digital user in chat applications; contact and messaging APIs link each person to a cyber account whose logs provide cues about availability, attention, and likely locations. Group channels supply shared context, forming a unified dialogue corpus that augments perception without additional fusion.

Smart-Lock Perceptor. Chassis robots equipped with smart

lockers use QR-code scanning and electronic lock-state feedback to confirm deposit and retrieval operations.

C. Simulation Environment

We build a text-based household simulator on top of TextWorld [38], extended to multi-room indoor scenes with manipulable objects and multiple embodied agents, mirroring domestic and office-like service tasks. Agents (robots or humans) can be issued with structured high-level commands (e.g., `navigate to <location>`, `pick up <object>`, `ask <human> for help`) and maintain an internal inventory of held items. To ensure interpretability, the simulator adopts deterministic symbolic actions rather than free-form language. All commands are parsed into templated feedback similar to ALFWorld [36], providing either confirmations for valid actions or standardized messages for unmet preconditions or incorrect contexts, enabling reliable multi-agent training and evaluation.

Robots are instantiated using the same configuration profiles as in Fig. 7. Each simulated robot model is governed by payload, workspace, and access constraints, and annotated with its available low-level skills. Attempts to execute actions beyond a robot’s capability deterministically fail, encouraging collaborative behaviors such as dynamic delegation and resource-aware planning. Multi-agent interactions, including object handovers, are explicitly modeled: a successful transfer requires an `offer` followed by a matching `take`, reflecting realistic turn-based coordination.

When a robot issues a help request (e.g., `ask human to lend <object>`), the simulator uses an LLM-driven conversational policy conditioned on task context. If the context specifies that a person is unavailable, the environment returns rejection or silence (e.g., “Sorry, I’m busy right now.”), capturing uncertainty in human-in-the-loop workflows. Although initial human information is provided as prior knowledge, availability may shift mid-task, requiring adaptive reasoning and replanning. This design provides a controlled but realistic foundation that supports consistent multi-agent evaluation while approximating key dynamics observed in the real-world system.

D. Baseline

We compare InteractGen with two types of complementary baselines with metrics shown in TABLE 2: **Generic Agent Reasoning Paradigms**—*ReAct* [17], *Reflexion* [18], and *Mobile-Agent-v2* [20] represent task-agnostic frameworks that interleave chain-of-thought reasoning with action execution or self-repair; although not tailored to embodied datasets, they supply strong multi-agent planning heuristics; **LLM-Powered Multi-Robot Controllers**—*CoELA* [23], *CaPo* [24], *RoCo* [25], *SMART-LLM* [22], *Lip-LLM* [26], *EMOS* [19], *LaMMA-P* [27], and *HMCF* [28], engineered for embodied tasks, exposing declarative skill APIs to dispatch heterogeneous robots thus serving as representative “LLM for multi-robot” baselines:

- **ReAct** [17]. Designed for single-agent tool use, ReAct interleaves chain-of-thought and actions but offers no primitives for robot selection or human queries. We instantiated *one* ReAct agent controlling all four robots sequentially and

TABLE 2
EVALUATION METRICS

SR	<i>Success Rate</i> — Percentage of fully completed instructions, allowing redundant steps
CR	<i>Completion Rate</i> — Proportion of expected interaction steps completed
RR	<i>Redundancy Rate</i> — Fraction of redundant actions among successful ones
RSR	<i>Robot Subtask Rate</i> — Proportion of subtasks completed by the robot out of all assignments
TOKEN	<i>Token Usage</i> — Average tokens used per successful execution, measured by the tokenizer
TIME	<i>Execution Time</i> — Total wall-clock time (in seconds) from task reception to completion

exposed a single “ask human” action. The agent seldom invoked this action, so ambiguities stayed unresolved; moreover, sequential control prevented true parallelism, bottlenecking multi-robot throughput.

- **Reflexion** [18]. Adds self-critique to ReAct. While this fixed minor execution slips, reflections occur *post hoc*; the agent still cannot clarify intent in advance, wasting steps before correcting. Parallel allocation issues persist because only one robot is active per reasoning step.
- **Mobile-Agent-v2** [20]. Originally maps GUI actions; we provide robot skills and augment GUI perception with text world summaries. The agent handles dynamic UI but assumes a single “device”, so concurrent robot scheduling degrades to serial execution; it barely requests human help and mis-allocates tasks requiring parallel execution.
- **CoELA** [23]. Accepts declarative skill APIs; we supplied four robot profiles and wrote wrappers that block skills a robot cannot perform. CoELA plans centrally but, lacking a clarification loop, silently proceeds when the goal is ambiguous, sometimes assigning unreachable subtasks. Without explicit concurrency control, task decomposition tends to under-utilise robots.
- **CaPo** [24]. Uses two LLM agents to co-optimize a plan; it can *in principle* ask the user, yet in practice the agents debate internally and rarely query the human. Divergent meta-plans stall progress under ambiguity; alternatively, both agents off-load the entire task to the human. Even with clear goals, its synchronous dialogue slows real-time execution and limits parallelism.
- **RoCo** [25]. Employs dialectic negotiation among robot-side LLMs, which rarely issue human queries—revealing a robot–robot rather than robot–human focus. Its distributed reasoning is latency-heavy: independent LLMs produce inconsistencies, leading to conflicting allocations.
- **SMART-LLM** [22]. Translates tasks into linear temporal logic for assignment. Although prompts encourage clarification, the agent seldom asks; ambiguity breaks the LTL translator. Dynamic human availability invalidates pre-solved plans, and it lacks a replanning trigger.
- **Lip-LLM** [26]. Couples an LLM with a linear-programming solver to optimize resource use, but no native clarification means wrong assumptions propagate.

TABLE 3

QUANTITATIVE COMPARISON ON OUR EVALUATION AGAINST BASELINES. ARROWS INDICATE DIRECTIONALITY OF BETTER PERFORMANCE. RATES ARE REPORTED AS PERCENTAGES (MEAN $\pm$ STD).

Method	SR [%] $\uparrow$		CR [%] $\uparrow$		RR [%] $\downarrow$		RSR [%] $\uparrow$		TOKEN ($\times 10^3$) $\downarrow$	TIME [s] $\downarrow$
	Sim	Real	Sim	Real	Sim	Real	Sim	Real	Real	Real
Generic Methods										
ReAct [17]	31 $\pm$ 4	13 $\pm$ 5	39 $\pm$ 3	21 $\pm$ 5	16 $\pm$ 3	21 $\pm$ 5	53 $\pm$ 3	50 $\pm$ 4	53.2 $\pm$ 4.1	2309 $\pm$ 119
Reflexion [18]	36 $\pm$ 3	19 $\pm$ 5	43 $\pm$ 2	27 $\pm$ 4	14 $\pm$ 3	19 $\pm$ 4	54 $\pm$ 3	50 $\pm$ 5	51.2 $\pm$ 4.1	2101 $\pm$ 111
Mobile-Agent-v2 [20]	54$\pm$3	43$\pm$3	54$\pm$5	49$\pm$3	8$\pm$3	13$\pm$4	65$\pm$4	54$\pm$5	43.2$\pm$3.0	1746$\pm$90
Multi-Robot Methods										
SMART-LLM [22]	47 $\pm$ 2	43 $\pm$ 4	53 $\pm$ 3	45 $\pm$ 4	17 $\pm$ 5	20 $\pm$ 1	59 $\pm$ 3	44 $\pm$ 3	49.4 $\pm$ 3.7	1996 $\pm$ 99
CoELA [23]	44 $\pm$ 1	29 $\pm$ 4	52 $\pm$ 3	41 $\pm$ 6	19 $\pm$ 2	23 $\pm$ 3	52 $\pm$ 2	42 $\pm$ 1	51.2 $\pm$ 4.1	2101 $\pm$ 111
Lip-LLM [26]	53 $\pm$ 2	45 $\pm$ 2	57 $\pm$ 4	48 $\pm$ 4	16 $\pm$ 4	19 $\pm$ 5	60 $\pm$ 5	49 $\pm$ 1	47.1 $\pm$ 3.4	1746 $\pm$ 93
RoCo [25]	50 $\pm$ 5	39 $\pm$ 4	53 $\pm$ 2	45 $\pm$ 1	17 $\pm$ 5	20 $\pm$ 3	55 $\pm$ 4	45 $\pm$ 3	48.7 $\pm$ 3.4	1823 $\pm$ 95
LaMMA-P [27]	56 $\pm$ 5	47 $\pm$ 7	62 $\pm$ 4	50 $\pm$ 6	12 $\pm$ 5	13 $\pm$ 6	62 $\pm$ 5	52 $\pm$ 5	44.1 $\pm$ 3.1	1597 $\pm$ 89
HMCF [28]	57 $\pm$ 5	51$\pm$1	66 $\pm$ 5	53 $\pm$ 4	8 $\pm$ 3	13 $\pm$ 5	62 $\pm$ 2	50 $\pm$ 1	43.7 $\pm$ 2.9	1504 $\pm$ 88
EMOS [19]	52 $\pm$ 5	48 $\pm$ 3	67 $\pm$ 4	54$\pm$2	12 $\pm$ 4	14 $\pm$ 3	65$\pm$2	53 $\pm$ 4	45.7 $\pm$ 3.1	1674 $\pm$ 89
CaPo [24]	62$\pm$4	50 $\pm$ 2	69$\pm$5	53 $\pm$ 4	7$\pm$3	11$\pm$5	63 $\pm$ 4	55$\pm$2	41.7$\pm$2.8	1367$\pm$88
InteractGen with Different Planning Foundation Models										
InteractGen (Planner on GPT-4o)	68 $\pm$ 3	57 $\pm$ 4	72 $\pm$ 4	63 $\pm$ 3	6 $\pm$ 3	5$\pm$2	66 $\pm$ 3	62 $\pm$ 2	36.1 $\pm$ 1.9	1137 $\pm$ 65
InteractGen (Planner on DS-R1)	74 $\pm$ 3	59 $\pm$ 5	75 $\pm$ 3	63 $\pm$ 3	6 $\pm$ 3	7 $\pm$ 3	78 $\pm$ 3	63 $\pm$ 4	34.7 $\pm$ 1.7	1061 $\pm$ 63
InteractGen (ToA Planner)	77$\pm$3	70$\pm$3	80$\pm$5	74$\pm$4	3$\pm$0	6 $\pm$ 1	81$\pm$3	68$\pm$5	31.3$\pm$1.3	970$\pm$56

The solver also ignores action duration, so simultaneous allocation overloads individual robots.

- **EMOS** [19]. Leverages structured “robot resumes” for reasoning but assumes fully specified task goals. It lacks language grounding to convert natural instructions into concrete predicates, often defaulting to arbitrary parameters. Once a plan is generated, it cannot be updated mid-execution—EMOS offers no support for human feedback or replanning under dynamic changes.
- **LaMMA-P** [27]. Integrates an LLM with a PDDL planner. We provided a handcrafted PDDL domain. Replanning is possible but slow—every environmental change triggers full PDDL reconstruction.
- **HMCF** [28]. Positions a human overseer who may intervene. We added a single “ask human for help” primitive; the LLM controller rarely used it, causing the human to intervene only *after* failures, increasing recovery cost. HMCF distributes one LLM per robot; concurrent reasoning raises bandwidth and consistency issues, making the overall coordination process a bottleneck.

To ensure a fair comparison, all baselines above are instantiated with the same foundation model, GPT-4o, following their original architectural design and prompting logic.

In InteractGen, the Perceiver and Assigner likewise uses GPT-4o, while the Verifier and Manager are implemented with DeepSeek-R1. Our Planner is trained in three stages:

STAGE 1: Supervised Fine-Tuning. Full-parameter SFT with DeepSpeed ZeRO-2 on eight A100 (bf16, 2k context). Per-GPU batch size 1 with 4 accumulation steps (effective batch size 32); learning rate 2×10^{-5} ; trained for 3 epochs.

STAGE 2: GRPO. Trained with ZeRO-2 under the same batch/precision settings for 5 epochs at 1×10^{-6} . Each query

samples 4 candidate plans to form GRPO groups. The KL penalty is removed to accelerate learning.

STAGE 3: Rejection Sampling Fine-Tuning. Full-parameter training resumes with ZeRO-3 on eight GPUs (bf16, 2k context). Per-GPU batch size 2 with 4 accumulation steps (effective batch size 64); learning rate 2×10^{-5} ; trained for 4 epochs. Rejection samples are generated using vLLM on a single A800 (temperature=1, top_p=0.7); invalid generations are discarded. The final planner is deployed under vLLM with temperature=0.2, top_p=0.95.

For the Manager’s clarification module, we use 700 ambiguous instructions from Section IV (600 train / 100 val) in the format *instruction* $\rightarrow$ *split JSON of task* $\rightarrow$ *masked key* $\rightarrow$ *clarification*. We fine-tune LLaMA-3.1-8B-Instruct using LoRA on a single RTX4090 (24 GB) for 10 epochs (bf16, 1024-token context). Inference uses deterministic decoding (temperature=0). Baseline models rely on prompt engineering to produce clarification when needed.

VII. RESULTS

In this section, we present the experimental results of InteractGen on multi-user, service tasks (see Fig. 8). We report comparative results against existing LLM-powered multi-agent frameworks and multi-robot systems, analyze the contribution of individual agent modules through ablations, and assess efficiency, usability, and user experience in real-world deployment.

A. Performance on Robot Coordination and Human Interaction

a) *InteractGen consistently outperforms all baselines.*: As shown in Table 3, *InteractGen* with the *ToA Planner*




Fig. 8. **Examples of multi-robot service execution in real-world open-use deployment.** InteractGen interprets natural-language requests from users and automatically assigns appropriate robot platforms to execute the required subtasks. Left: a mobile robot is instructed to fetch a drink and escort it to the user. Middle: a mobile manipulator retrieves a bottle from the tea room and delivers it to the requester. Right: a quadruped robot picks up mineral water from a designated location and brings it to the user, demonstrating adaptive task allocation across heterogeneous robots and dynamic human-robot environments.




Fig. 9. **Real-world performance comparison across six evaluation metrics.** Box plots report (a) Success Rate, (b) Completion Rate, (c) Redundant Rate, (d) Robot Subtask Rate, (e) Token Usage, and (f) Execution Time for five representative methods: Reflexion, HMCf, EMOS, CaPo, and InteractGen (ToA). InteractGen achieves the highest SR, CR, and RSR while maintaining the lowest redundancy and resource cost, as reflected by reduced token usage and shorter execution time. Baseline methods exhibit lower median performance and larger variance across metrics, highlighting InteractGen’s superior consistency, efficiency, and robustness in real-world human-robot interaction tasks.

achieves the best overall performance. In simulation, it reaches the highest Success Rate (SR = **0.77**) and Completion Rate (CR = **0.80**), outperforming the strongest baseline, *CaPo* (SR = 0.62, CR = 0.69), by over 15%. Redundancy Rate (RR) drops to **0.03**, and Robot Subtask Rate (RSR) rises to **0.81**. Similar trends are observed in the real world (SR = **0.70**, CR = **0.74**). These gains stem from ToA’s modular and interpretable design, where planning, assignment, and validation are explicitly structured, tightly coordinated, and grounded in fine-grained task dependencies. While planning with GPT-4o lowers RR (0.05 vs. ToA’s 0.06), its flatter reasoning often misses dependencies, whereas ToA achieves better global consistency through explicit, fine-grained subtask structuring. *InteractGen* consistently maintains a high success rate while

requiring remarkably minimal human assistance, even as the complexity of human-object interactions substantially increases (see Fig. 9 and Fig. 10 for more intuitive visual analyses).

b) Structured ToA Planner improves robustness beyond general-purpose LLMs. ToA’s advantage also comes from its dedicated training pipeline for long-horizon, multi-agent planning. Unlike general-purpose LLMs like GPT-4o, ToA encodes structured reasoning patterns that improve task decomposition and allocation. This leads to fewer planning errors and better resource utilization, as reflected in its top RSR (**0.81**) and lowest RR (**0.03**). It even outperforms the deep-thinking model DeepSeek-R1, demonstrating that careful structural design can surpass raw reasoning depth alone. These results confirm that *Thought-of-Action* planning yields more reliable




Fig. 10. **Performance over normalized task steps.** Task steps are normalized by the number of required human interactions and object-related operations in each task, providing a unified measure of progress across scenarios. (a) Success rate, (b) completion rate, and (c) redundancy rate as tasks advance. Our ToA-based method sustains higher success and completion rates and consistently lower redundancy throughout the entire execution horizon, whereas baseline methods degrade more rapidly as the number of required interactions increases.




Fig. 11. **Action distribution across different interaction stages.** Box plots show the number of robot actions (left) and human actions (right) generated by each method at three intervals: (a) steps 3-7, (b) steps 8-12, and (c) steps 13-15. InteractGen (ToA) consistently produces fewer and more focused robot and human action requests compared with CaPo, EMOS, HMCf, and Reflexion, indicating more efficient reasoning and reduced unnecessary agent calls. As tasks progress to later stages, baseline methods exhibit increasing variance and action redundancy, whereas InteractGen maintains stable and concise action generation.

behavior than generic LLM prompting, offering clearer task decomposition and stronger execution consistency (see Fig. 11).

c) Decentralized methods exhibit lower reliability and weak human-robot coordination.: Baselines like *SMART-LLM*, *HMCf*, and *CaPo* rely on decentralized multi-LLM setups, where each robot runs its own planner. This often leads to redundant or conflicting behaviors, especially under partial observability or limited communication. For instance, *HMCf* reports $SR = 0.57$ and $RR = 0.08$ in simulation, while *CaPo* shows $CR = 0.69$ but suffers from poor human subtask allocation ($RSR = 0.63$). Without a dedicated centralized plan backbone, agents frequently act sequentially or largely rely on human agent, reducing overall task success and parallelism.

B. Token and Time Efficiency Analysis

Despite coordinating five specialized agents, *InteractGen* achieves the highest token efficiency among all compared systems. It uses only 31.3k tokens per successful run—about 40% fewer than single-agent baselines such as Reflexion (51.2k) and ReAct (53.2k), and 25–35% fewer than decentralized controllers including CaPo (41.7k), HMCf (43.7k), and Lip-LLM (47.1k). This efficiency stems from three design choices:

(i) the Perceiver emits only incremental state updates rather than replaying full histories; (ii) the Planner produces typed, dependency-aware Thought-of-Action graphs rather than verbose free-form CoT; and (iii) all agents communicate via concise, role-specific prompts that avoid repeated grounding and unnecessary self-dialogue. These mechanisms curb token growth while preserving interpretability, enabling graceful scaling with task horizon and team size.

These design principles also yield strong time efficiency. *InteractGen* completes tasks in 970 s on average, outperforming the fastest baseline (CaPo, 1367 s) by 29% and ReAct (2309 s) by over 55%. Centralized planning and assignment allow independent subtasks to run in parallel, eliminating the pairwise negotiation overhead that causes near-quadratic message growth in decentralized methods. The Validator prevents time-consuming execution of infeasible plans, and the Manager resolves ambiguities early, reducing downstream failure recovery. Ablations reinforce these observations: removing the Perceiver inflates token usage, removing the Validator prolongs fail–retry cascades, and removing the Manager introduces redundant clarifications. These results show that *InteractGen*’s modular pipeline is not only more robust than monolithic reasoning but




Fig. 12. **Ablation analysis of InteractGen.** (a–b) 3D visualizations of Success Rate (SR), Completion Rate (CR), and Redundant Rate (RR) under module ablations in both simulation and real-world settings, highlighting the performance degradation caused by removing key agents. (c–e) Corresponding bar plots of SR, CR, and RR, showing consistent trends across environments and revealing each agent function as a critical component for reliable task execution.




Fig. 13. **User survey results from a multi-month real-world deployment.** Left: the distribution of embodied agents involved in tasks shows that InteractGen effectively coordinates heterogeneous robots while keeping human participation lightweight. Right: aggregated user feedback indicates clear reductions in perceived labor, high satisfaction with system behavior, appropriate levels of human-agent requests, and noticeable improvements in overall productivity. These findings demonstrate InteractGen’s strong usability and its balanced integration of automation and human collaboration in practical service environments.

also computationally lean, offering both high success rates and practical efficiency for real-world human-robot collaboration.


C. Ablation Studies on Contributions of Individual Agents

We assess the functional role of each agent by removing one module at a time and replacing its behavior with a minimal prompt attached to the Assigner. This forces the system to operate without the corresponding capability and exposes its necessity. As shown in Fig. 12, performance degrades

consistently under every ablation, demonstrating that the five-agent decomposition is structurally indispensable.

The Planner is the most critical component: removing structured task decomposition causes the success rate (SR) to fall from 0.68 to 0.45, while repetition rate (RR) nearly doubles, reflecting incoherent and looping plans typical of monolithic reasoning. Removing the Perceiver also leads to large drops (SR: 0.58), as the system loses incremental world-state updates and fails to react to dynamic changes such as user relocation. Ablating the Validator reduces SR to 0.60 because unsafe or

Proactive 1: Alan wants Bob to print a document and Carl to sign it.



Proactive 2: Eric spilled water on the desk and wants some tissue.




Fig. 14. **Proactive1:** The system adapts to human-context changes before failure occurs. Originally planning to deliver a document to Carl's office, it detects inconsistencies—Service Bot1 not meeting Carl, unused smart lock, and a group message showing Carl in Room 401E. The Manager infers Carl has moved, triggering re-planning; the Planner redirects the robot to 401E, where Carl signs the document. This case shows proactive tracking of human movement and assumption revision. **Proactive2:** Eric requests tissue after spilling water. The system identifies Bob as a source and dispatches Service Bot1, but Bob reports he is off-site and left the tissue on his desk. The system then uses RealSense to confirm the tissue's location and dispatches Service Bot2 (with a manipulator) to grasp and deliver it. Throughout, *InteractGen* updates Eric, adapts plans from feedback, and performs the task efficiently—demonstrating robust sensing, coordination, and human-aware planning.

infeasible actions bypass pre-execution checks, resulting in avoidable failures. Removing the Manager limits adaptability (SR: 0.61); without reflection and recovery mechanisms, the system cannot correct downstream errors.

These ablations verify that each LLM agent contributes a distinct and irreplaceable capability. The resulting archi-

ture is modular, interpretable, and far more stable than a monolithic chain-of-thought pipeline. Even when compared against stronger proprietary baselines such as DeepSeek-R1, the fully assembled *InteractGen* system achieves superior robustness and execution quality, underscoring the practical necessity of modular specialization. We showcase two cases

in Fig. 14, demonstrating how *InteractGen* executes real-world tasks through active adaptation and proactive querying, leveraging self-reflection to coordinate robots and humans.

D. Real-World Performance Assessment

We conducted a three-month open-use study (Fig 13) in a shared academic office equipped with our heterogeneous robot fleet and human participants accessible via a chat interface. Users issued natural-language requests which *InteractGen* parsed and executed through its multi-agent pipeline. Subtasks were allocated dynamically: mobile chassis robots handled secure deliveries, the quadruped manipulator managed unstructured navigation and pickups, and humans were involved only for socially sensitive or dexterous operations. Across 873 valid survey responses, users reported a 73% satisfaction rate, strong perceived labor reduction (4.8/5), and high comfort with system-issued human requests (4.2/5). Furthermore, 81% agreed that *InteractGen* improved their productivity, indicating substantial real-world impact beyond technical feasibility.

Operational statistics further show how the system coordinated mixed teams effectively. Humans performed 31% of subtasks—mainly fine-motor or judgment-heavy actions—while mobile robots completed 42% and manipulators 22%, dominated by secure deliveries and simple fetch-and-carry tasks. Other platforms were used less frequently, matching the natural distribution of office demands. This division of labor highlights how *InteractGen* balances autonomy with inclusivity: robots absorb routine and mobility-intensive work, while humans are engaged only when their unique strengths are essential.

Qualitative feedback emphasized reduced physical strain, lower cognitive load in coordinating shared resources, and improved fairness in delegation, as system-generated requests were perceived as neutral and equitable. These findings demonstrate that *InteractGen* is not only efficient and reliable, but also socially adaptive—reinforcing its suitability as a deployable, human-centered assistant in everyday work environments.

VIII. DISCUSSION

A. The Potential for LLM Agents in Embodied Tasks

LLMs have progressed from static predictors to interactive, tool-using multi-agent systems. Yet most advances remain confined to digital settings, leaving open whether such architectures can meaningfully support embodied agents operating in dynamic, human-centered service workflows.

Our study shows that these principles transfer naturally to embodied contexts. By organizing *InteractGen* as a five-agent pipeline for perception, planning, decision, validation, and reflection, we enabled heterogeneous robots and humans to coordinate effectively in office-service routines. This modular decomposition improves interpretability and, crucially, treats humans as deployable agents rather than passive overseers. Across extensive experiments, this architecture consistently surpassed related baselines in success rate, adaptability, latency, and token efficiency, demonstrating that the collaborative strengths of digital multi-agent frameworks yield concrete performance gains in embodied workflows (see Fig. 15).




Fig. 15. **Method rankings across six metrics.** *InteractGen* variants remain near the top with minimal fluctuation, demonstrating stable advantages across all dimensions. In contrast, existing baselines vary significantly across metrics, revealing trade-offs or weaknesses.

A key advantage lies in robustness through structured delegation. In deployment, when robots encountered limitations—such as fragile object handling or user locations that were not precisely known—the reflection module dynamically reassigned subtasks to available humans or re-routed the workflow to maintain continuity. This pattern highlights a broader capability: LLM-based agent architectures can integrate diverse skills, embed human agency into the control loop, and sustain resilience in dynamic environments. We contend that such design principles represent a meaningful step toward embodied systems that are both technically competent and socially adaptive, narrowing the gap between automation and real-world human–robot collaboration.

B. Foundation Models for Constructing Specialized Agents

To enable robust human–robot mixed-team coordination, we design a comprehensive multi-agent reasoning pipeline spanning perception, planning, decision, validation, and reflection. Each stage is mapped directly to a specialized LLM agent, mirroring the cognitive workflow of a human assistant: extracting intent, perceiving the environment, generating a roadmap, selecting immediate actions, validating feasibility, requesting human help when needed, and reflecting on outcomes. This decomposition provides modularity and interpretability, while also establishing a foundation for treating humans as integral, deployable agents rather than passive overseers.

For instantiating these roles, foundation models such as GPT [40], DeepSeek [41], and Gemini [42] series offer state-of-the-art generalization and multimodal capacity. However, when domain-specific data is available, smaller open-weight models—including LLaMA3.1-8B-Instruct, Qwen3-8B, and Mistral-8B-Instruct—can be fine-tuned for planning and decision stages, yielding strong performance at lower cost. In contrast, perception, validation, and reflection stages typically benefit from large-scale pretrained models due to their

richer contextual grounding. Careful prompt specialization [43] and targeted few-shot examples [44] further enhance reasoning quality and consistency across agents.

Our experiments validate this design. Table 4 shows that fine-tuned open-weight models can already outperform proprietary baselines after STAGE 2, with Qwen3-8B achieving the strongest gains. Its success rate (SR) and completion rate (CR) improved from 0.25/0.53 at STAGE 1 to **0.77/0.80** at STAGE 3, also yielding the lowest replan rate (RR=0.03). LLaMA3.1-8B and Mistral-8B likewise approached Claude-3.7-Sonnet performance by STAGE 3. These results highlight two key insights: (i) multi-stage fine-tuning with role specialization is crucial for embodied multi-agent systems, and (ii) well-tuned open-weight models can rival or surpass proprietary alternatives, making deployment more scalable and cost-effective.

C. Scalable Coordination and Generalization in InteractGen

Single-agent paradigms such as ReAct [17] and Reflexion [18] entangle perception, planning, and reflection into a monolithic chain of thought, often yielding incomplete reasoning and inefficient execution. Decentralized multi-robot systems like CoELA [23] and RoCo [25] distribute reasoning but suffer from dialogue-induced latency that scales poorly with agent count. *InteractGen* provides a middle ground through a modular five-agent pipeline in which each agent performs independent reasoning and exchanges only concise, role-specific messages. This design delivers both scalability and interpretability, enabling efficient execution of complex service workflows in multi-human, multi-robot environments.

Modular specialization brings robustness and adaptive recovery. The Validator functions as a safety gate for filtering infeasible or unsafe outputs, while the Perceiver curbs context inflation by maintaining incremental world states rather than replaying full histories. Removing these modules produces sharp performance drops, underscoring their necessity. Combined with real-time perception and LLM-as-judge reflection, *InteractGen* supports a reliable *act* $\rightarrow$ *fail* $\rightarrow$ *reflect* $\rightarrow$ *replan* loop that detects anomalies, triggers replanning, and reassigns tasks—including invoking human assistance when robots reach intrinsic limits. Fragile or thin objects (e.g., documents) are automatically routed to humans, and repeated robotic failures on unfamiliar items lead to empirical adjustments in future task allocation. This reflective cycle allows the system to evolve beyond static rules and embed humans as integral agents in the collaboration loop.

Both simulation and real-world evaluations confirm the scalability and generalization enabled by this architecture. During the evaluation process, *InteractGen* more than doubled the success and completion rates of ReAct-style baselines. The three-stage ToA planner (SFT $\rightarrow$ GRPO $\rightarrow$ Rejection Sampling) outperformed significantly larger proprietary models while remaining more cost-efficient. During three months of open deployment, the system handled unscripted tasks such as document signing and food delivery, with surveys reporting 73% user satisfaction and 81% productivity improvement across 873 responses. These results establish *InteractGen* as a robust,

TABLE 4
EVALUATION IN THE SIMULATION ENVIRONMENT OF DIFFERENT MODELS AS PLANNER FOUNDATION MODELS. THE QWEN3-8B MODEL, AFTER THREE-STAGE TRAINING, ACHIEVES THE BEST PERFORMANCE AND IS USED AS THE ToA PLANNER IN INTERACTGEN.

Model		SR↑	CR↑	RR↓	RSR↑
Proprietary					
	Claude-3.7-Sonnet-thinking	75	77	7	71
	GPT-o3-mini	75	79	7	73
	GPT-4o	68	72	6	66
	GPT-o1-mini	70	73	5	68
Open-weight					
	DeepSeek-R1	74	75	6	78
	LLaMA3.1-8B-Instruct	51	60	23	53
	Ministral-8B-Instruct	54	63	20	59
	Qwen3-8B	60	66	14	64
Finetuned					
STAGE 1	LLaMA3.1-8B-Instruct	60	66	18	59
	Ministral-8B-Instruct	63	67	16	59
	Qwen3-8B	68	70	10	67
STAGE 2	LLaMA3.1-8B-Instruct	66	70	8	63
	Ministral-8B-Instruct	69	73	8	68
	Qwen3-8B	73	75	5	74
STAGE 3	LLaMA3.1-8B-Instruct	71	74	7	73
	Ministral-8B-Instruct	73	77	6	76
	Qwen3-8B	77	80	3	81

scalable, and practically deployable framework for dynamic human-robot collaboration.

D. Social Utility and Real-World Impact

Beyond technical performance, *InteractGen* demonstrates clear real-world utility. Unlike systems restricted to scripted benchmarks or controlled labs, it operates in dynamic multi-user service workflows. Its structured Thought-of-Action planning, reflective validation, and proactive perception enable robust execution of long-horizon tasks—such as printing, signing, and document delivery—and consistently outperform decentralized baselines in both simulation and deployment.

This robustness translated directly into practice. In a three-month open deployment in a shared academic office, *InteractGen* reliably handled natural user instructions, coordinated heterogeneous robots, and managed shared resources with minimal intervention. Users reported reduced physical effort and cognitive load, smoother collaboration, and improved inclusivity, as the system’s neutral request style mitigated interpersonal friction and encouraged mutual assistance.

Figure 16 illustrates a representative example. We integrate DBNet and ConvNeXtViT-document for OCR, GroundingDINO for icon grounding, and Qwen-VL-Plus for GUI captioning to enable multimodal perception and GUI-level operation via simulated tapping. This allows workflows such as one-shot food ordering: the system navigates the interface, coordinates robotic delivery, and only prompts users for payment. The deployment also included wheeled and quadruped robots (Fig. 8), requiring minimal adaptation, demonstrating the generality of the overall architecture.




Fig. 16. **A hybrid digital-physical workflow.** Given the instruction “Please order me a takeaway of beef noodles,” InteractGen coordinates digital actions with real-world robot execution. *Cyber Display*: InteractGen searches for the dish, confirms the order, forwards the payment page to the user, completes payment, informs the delivery person through the app, and sends status updates and QR codes to the user. *Real-world Display*: the mobile robot navigates to the pickup point, receives the takeaway from the delivery person, and returns to deliver it to the requester. The sequence illustrates how InteractGen seamlessly integrates online service operations with physical robot actions to accomplish multi-step tasks across digital and real environments.

Together, these results highlight both the technical and social significance of *InteractGen*. Technically, it provides a scalable and interpretable multi-agent reasoning framework that supports reliable, adaptive execution in dynamic environments. Socially, it acts as a coordination mediator that distributes tasks more fairly, reduces collaboration burden, and improves accessibility for users with limited mobility or time. This dual impact positions *InteractGen* not merely as an automation system, but as a step toward sustainable, human-centered teamwork suitable for workplaces, assistive settings, and other socially embedded domains.

IX. CONCLUSION

This work re-examines the role of foundation models in robotics through the lens of real-world service workflows. Rather than asking whether a single Vision-Language Model or Vision-Language-Action policy can serve as a universal robot brain, we showed that the core limitation is architectural: monolithic models are fundamentally misaligned with the distributed, evolving, and dependency-heavy nature of human-populated environments. Building on this perspective, we introduced *InteractGen*, an LLM-powered multi-agent framework that treats foundation models as regulated components within a structured control system rather than as end-to-end embodied policies. By decomposing robot intelligence into specialized agents for perception, planning, assignment, validation, reflection, and human delegation, *InteractGen* demonstrates that reliable autonomy emerges from coordinated roles, shared memory, and

explicit feasibility checks. The framework delivers consistent gains across simulation and real-world deployment, including a three-month open-use study in a shared office, where it executed long-horizon tasks with high reliability, adapted fluidly to environmental changes, and collaborated with humans in natural and socially appropriate ways. Even low-cost mobile robots exhibited markedly elevated capability when embedded in *InteractGen*’s human-aware multi-agent loop, handling tasks that far exceed the reach of standalone foundation models.

At the same time, our study has limitations that point to rich future directions. *InteractGen* is currently constrained by hand-engineered agent roles and does not explicitly demonstrate emergent abilities. Extending the framework to learn skills, role decompositions, and inter-agent protocols directly from data would further improve scalability and broaden domain coverage. Our experiments focus on domestic-office workflows with moderate team sizes; deploying similar architectures in safety-critical domains such as healthcare, eldercare, or public spaces will require tighter integration with formal safety guarantees, richer uncertainty modeling, and stronger robustness to adversarial or strategically behaving humans. Moreover, while our results highlight the social utility of treating humans as active agents, a deeper understanding of long-term trust, fairness, and workload distribution in human-robot-agent collectives remains an open challenge.

We view *InteractGen* not as a final answer, but as a compelling existence proof that foundation models, when embedded within carefully structured multi-agent control systems,

can substantially advance service robotics toward intentional, human-centered autonomy with measurable real-world impact. More broadly, the framework illustrates how robots and LLM-driven agents can move beyond reactive execution toward collaborative problem solving with people—supporting clearer communication, more adaptive task sharing, and safer, more trustworthy human–robot interaction. We believe such systems mark a step toward assistive technologies that amplify human capability, reduce cognitive and physical burden, and enable more equitable access to real-world robotic assistance.

REFERENCES

- [1] X. Zhang *et al.*, “Llm-grop: Visually grounded robot task and motion planning with large language models,” *The International Journal of Robotics Research*, 2024.
- [2] Y. Xu *et al.*, ““Set It Up”: Functional object arrangement with compositional generative models,” *The International Journal of Robotics Research*, 2024.
- [3] D. Lee, S. Joo, K. Lee, and B. Kim, “Prime the search: Using large language models for guiding geometric task and motion planning by warm-starting tree search,” *The International Journal of Robotics Research*, 2024.
- [4] R. Firoozi *et al.*, “Foundation models in robotics: Applications, challenges, and the future,” *The International Journal of Robotics Research*, 2025.
- [5] G. Luo, G. Yang *et al.*, “Visual embodied brain: Let multimodal large language models see, think, and control in spaces,” 2025. [Online]. Available: <https://arxiv.org/abs/2506.00123>
- [6] Y. Ji, H. Tan *et al.*, “Robobrain: A unified brain model for robotic manipulation from abstract to concrete,” 2025. [Online]. Available: <https://arxiv.org/abs/2502.21257>
- [7] X. Hao, L. Zhou, Z. Huang *et al.*, “Mimo-embodied: X-embodied foundation model technical report,” 2025. [Online]. Available: <https://arxiv.org/abs/2511.16518>
- [8] M. Ahn, A. Brohan, N. Brown, Y. Chebotar *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” in *Proceedings of the Conference on Robot Learning*. PMLR, 2022. [Online]. Available: <https://arxiv.org/abs/2204.01691>
- [9] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, K. Choromanski *et al.*, “Rt-2: Vision-language-action models transfer web knowledge to robotic control,” *Proceedings of the Conference on Robot Learning*, vol. 229, pp. 2165–2183, 2023. [Online]. Available: <https://proceedings.mlr.press/v229/zitkovich23a.html>
- [10] M. J. Kim, K. Pertsch *et al.*, “Openvla: An open-source vision-language-action model,” *Proceedings of the Conference on Robot Learning*, vol. 270, pp. 2679–2713, 2025. [Online]. Available: <https://proceedings.mlr.press/v305/black25a.html>
- [11] K. Black, N. Brown *et al.*, “ π_0 : A vision-language-action flow model for general robot control,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.24164>
- [12] Q. Li, Y. Liang *et al.*, “Cogact: A foundational vision-language-action model for synergizing cognition and action in robotic manipulation,” 2024. [Online]. Available: <https://arxiv.org/abs/2411.19650>
- [13] P. Intelligence *et al.*, “ $\pi_{0.5}$: a vision-language-action model with open-world generalization,” in *Proceedings of the Conference on Robot Learning*. Proceedings of Machine Learning Research (PMLR), 2025. [Online]. Available: <https://proceedings.mlr.press/v305/black25a.html>
- [14] Q. Wu, G. Bansal *et al.*, “Autogen: Enabling next-gen llm applications via multi-agent conversation,” 2023. [Online]. Available: <https://arxiv.org/abs/2308.08155>
- [15] S. Hong, M. Zhuge *et al.*, “Metagpt: Meta programming for a multi-agent collaborative framework,” in *International Conference on Learning Representations*, 2023. [Online]. Available: <https://api.semanticscholar.org/CorpusID:265301950>
- [16] Z. Chen, H. Zhang, Y. Liang, Q. Li, Y. Li, and J. Tang, “Agentverse: Facilitating multi-agent collaboration and exploration with large language models,” *Proceedings of the International Conference on Learning Representations*, 2024. [Online]. Available: <https://arxiv.org/abs/2308.10848>
- [17] S. Yao, J. Zhao *et al.*, “React: Synergizing reasoning and acting in language models,” in *Proceedings of the 11th International Conference on Learning Representations*. Kigali, Rwanda: International Conference on Learning Representations, 2023. [Online]. Available: <https://arxiv.org/abs/2210.03629>
- [18] N. Shinn, F. Cassano *et al.*, “Reflexion: Language agents with verbal reinforcement learning,” in *The Twelfth International Conference on Learning Representations*, 2024, international Conference on Learning Representations (ICLR). [Online]. Available: <https://arxiv.org/abs/2303.11366>
- [19] J. Chen *et al.*, “Emos: Embodiment-aware heterogeneous multi-robot operating system with llm agents,” 2025. [Online]. Available: <https://arxiv.org/abs/2410.22662>
- [20] J. Wang *et al.*, “Mobile-agent-v2: Mobile device operation assistant with effective navigation via multi-agent collaboration,” 2024. [Online]. Available: <https://arxiv.org/abs/2406.01014>
- [21] A. Z. Ren, A. Dixit *et al.*, “Robots that ask for help: Uncertainty alignment for large language model planners,” in *Proceedings of the 2023 Conference on Robot Learning*, 2023. [Online]. Available: <https://arxiv.org/abs/2307.01928>
- [22] S. S. Kannan, V. L. N. Venkatesh, and B.-C. Min, “Smart-llm: Smart multi-agent robot task planning using large language models,” in *Proceedings of the IEEE International Conference on Robotics and Automation*, 2024. [Online]. Available: <https://arxiv.org/abs/2309.10062>
- [23] H. Zhang, W. Du *et al.*, “Building cooperative embodied agents modularly with large language models,” in *Proceedings of the International Conference on Learning Representations*, 2024. [Online]. Available: <https://arxiv.org/abs/2307.02485>
- [24] J. Liu, P. Zhou, Y. Du, A.-H. Tan, C. G. M. Snoek, J.-J. Sonke, and E. Gavves, “Capo: Cooperative plan optimization for efficient embodied multi-agent cooperation,” 2025. [Online]. Available: <https://arxiv.org/abs/2411.04679>
- [25] Z. Mandi, S. Jain, and S. Song, “Roco: Dialectic multi-robot collaboration with large language models,” 2023. [Online]. Available: <https://arxiv.org/abs/2307.04738>
- [26] K. Obata *et al.*, “Lip-llm: Integrating linear programming and dependency graph with large language models for multi-robot task planning,” 2024. [Online]. Available: <https://arxiv.org/abs/2410.21040>
- [27] X. Zhang, H. Qin, F. Wang, Y. Dong, and J. Li, “Lamma-p: Generalizable multi-agent long-horizon task allocation and planning with llm-driven pddl planner,” 2025. [Online]. Available: <https://arxiv.org/abs/2409.20560>
- [28] Z. Li *et al.*, “Hmcf: A human-in-the-loop multi-robot collaboration framework based on large language models,” 2025. [Online]. Available: <https://arxiv.org/abs/2505.00820>
- [29] A. Radford, J. W. Kim *et al.*, “Learning transferable visual models from natural language supervision,” 2021. [Online]. Available: <https://arxiv.org/abs/2103.00020>
- [30] L. Beyer *et al.*, “Paligemma: A versatile 3b vision-language model for transfer,” 2024. [Online]. Available: <https://arxiv.org/abs/2407.07726>
- [31] L. Vig and J. Adams, “Multi-robot coalition formation,” *IEEE Transactions on Robotics*, vol. 22, no. 4, pp. 637–649, 2006.
- [32] S. S. O. Venkata, R. Parasuraman, and R. Pidaparti, “Kt-bt: A framework for knowledge transfer through behavior trees in multirobot systems,” *IEEE Transactions on Robotics*, vol. 39, no. 5, pp. 4114–4130, 2023.
- [33] R. Banerjee, R. K. Jenamani *et al.*, “To ask or not to ask: Human-in-the-loop contextual bandits with applications in robot-assisted feeding,” 2025. [Online]. Available: <https://arxiv.org/abs/2405.06908>
- [34] M. M. Veloso, “Symbiotic-autonomous service robots for user-requested tasks in a multi-floor building,” <https://doi.org/10.1184/R1/6610058.v1>, 2018.
- [35] N. Sun *et al.*, “Assistantx: An llm-powered proactive assistant in collaborative human-populated environments,” in *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025, pp. 3352–3359.
- [36] M. Shridhar, X. Yuan *et al.*, “Alfworld: Aligning text and embodied environments for interactive learning,” in *The Ninth International Conference on Learning Representations*, 2021, international Conference on Learning Representations (ICLR). [Online]. Available: <https://arxiv.org/abs/2010.03768>
- [37] S. Srivastava, C. Li *et al.*, “Behavior: Benchmark for everyday household activities in virtual, interactive, and ecological environments,” 2021. [Online]. Available: <https://arxiv.org/abs/2108.03332>
- [38] M.-A. Côté *et al.*, “Textworld: A learning environment for text-based games,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019. [Online]. Available: <https://arxiv.org/abs/1806.11532>
- [39] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, “Learning fine-grained bimanual manipulation with low-cost hardware,” 2023. [Online]. Available: <https://arxiv.org/abs/2304.13705>
- [40] OpenAI *et al.*, “Gpt-4 technical report,” 2024. [Online]. Available: <https://arxiv.org/abs/2303.08774>
- [41] DeepSeek-AI *et al.*, “Deepseek-v3 technical report,” 2025. [Online]. Available: <https://arxiv.org/abs/2412.19437>

- [42] G. Team, R. Anil *et al.*, “Gemini: A family of highly capable multimodal models,” 2025. [Online]. Available: <https://arxiv.org/abs/2312.11805>
- [43] Q. Ye, M. Axmed, R. Pryzant, and F. Khani, “Prompt engineering a prompt engineer,” 2024. [Online]. Available: <https://arxiv.org/abs/2311.05661>
- [44] J. Coda-Forno *et al.*, “Meta-in-context learning in large language models,” 2023. [Online]. Available: <https://arxiv.org/abs/2305.12907>