

Leveraging Adaptive Group Negotiation for Heterogeneous Multi-Robot Collaboration with Large Language Models

Anonymous ACL submission

Abstract

Multi-robot collaboration tasks often require heterogeneous robots to work together over long horizons under spatial constraints and environmental uncertainties. Although Large Language Models (LLMs) excel at reasoning and planning, their potential for coordinated control has not been fully explored. Inspired by human teamwork, we present CLiMRS (Cooperative Large-Language-Model-Driven Heterogeneous Multi-Robot System), an adaptive group negotiation framework among LLMs for multi-robot collaboration. This framework pairs each robot with an LLM agent and dynamically forms subgroups through a general proposal planner. Within each subgroup, a subgroup manager leads perception-driven multi-LLM discussions to get commands for actions. Feedback is provided by both robot execution outcomes and environment changes. This grouping–planning–execution–feedback loop enables efficient planning and robust execution. To evaluate these capabilities, we introduce CLiMBench, a heterogeneous multi-robot benchmark of challenging assembly tasks. Our experiments show that CLiMRS surpasses the best baseline, achieving over 40% higher efficiency on complex tasks without sacrificing success on simpler ones. Our results demonstrate that leveraging human-inspired group formation and negotiation principles markedly enhances the efficiency of heterogeneous multi-robot collaboration. Our code is available [here](#).

1 Introduction

Addressing real-world, everyday tasks often requires collaboration to efficiently handle long-horizon planning and complex perception. However, developing embodied agents for multi-robot tasks remains an open challenge. Inspired by human teamwork, incorporating human teaming principles into multi-robot collaboration, where groups of robotic agents coordinate planning and perception through shared observations and information,

offers a promising yet challenging path to improving efficiency and robustness (Zhang et al., 2024b).

Meanwhile, large language models (LLMs) have demonstrated strong capabilities across a range of tasks, including question answering (Rein et al., 2024), code generation (Jain et al., 2024), and logical reasoning (Plaat et al., 2024). Recent work has integrated LLMs into robotic planning (Song et al., 2023; Zhang et al., 2024a; Mower et al., 2024; Salimpour et al., 2025; Liang et al., 2025), with several studies extending these approaches to multi-robot collaboration (Zhang et al., 2024b; Mandi et al., 2024; Liu et al., 2025). However, previous work has mainly focused on homogeneous agents (Liu et al., 2024a), limiting the diversity of capabilities that can be exhibited during collaboration. Moreover, existing work on heterogeneous teams often assumes idealized operating conditions (Liu et al., 2025), overlooking cumulative errors over long horizons, underestimating communication costs, and overestimating cooperative efficiency. As a result, despite the promise of LLM-driven multi-robot collaboration, significant gaps remain under heterogeneous agents, long-horizon objectives, and real-world constraints.

To address these limitations, we propose CLiMRS (Cooperative Large-Language-Model-Driven Heterogeneous Multi-Robot System), an adaptive group negotiation framework among multiple LLMs that enables multi-robot collaboration. CLiMRS orchestrates heterogeneous robots via dynamic subgroup formation and cooperative planning, enabling robust long-horizon collaboration in uncertain environments. Within this framework, each robot is paired with an independent LLM agent that communicates with its peers to accomplish complex, long-horizon tasks. To enhance collaborative effectiveness, the system leverages the broad world knowledge of LLMs and explicitly models inter-agent dependencies through a grouping–planning–execution–feedback loop.

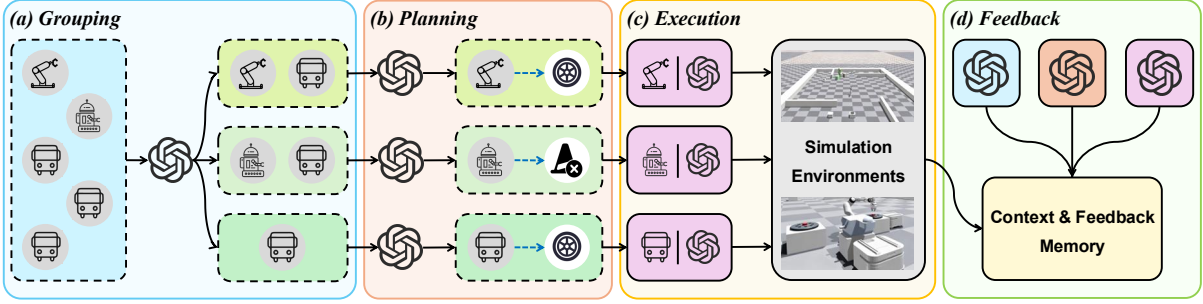


Figure 1: **Overview.** We present CLiMRS, an adaptive group negotiation framework among multiple LLMs that enables multi-robot collaboration through a grouping–planning–execution–feedback loop, and CLiMBench, a heterogeneous multi-robot collaboration benchmark in simulation with challenging long-horizon assembly tasks.

To systematically evaluate the applicability of CLiMRS in challenging scenarios where heterogeneous robots must cope with unpredictable execution errors, we introduce **CLiMBench**, a benchmark for heterogeneous multi-robot collaboration. CLiMBench is built around long-horizon assembly tasks that require robots to jointly perform object search, navigation, transportation, and assembly under partial observability. It features five robotic devices across three types of heterogeneous robots. Tasks of varying difficulty simulate material-handling and assembly processes with diverse skill usage, designed to test the planning and perception capabilities of LLM-based frameworks.

We evaluated our framework in CLiMBench and another heterogeneous robot collaboration benchmark (Liu et al., 2025). Our experiments show that CLiMRS outperforms the best baseline, increasing success rates and improving efficiency on complex tasks while maintaining high success on simpler ones. These results demonstrate that incorporating human-inspired dynamic subgroup formation and negotiation principles substantially enhances the efficiency of heterogeneous multi-robot collaboration. To summarize, our main contributions are:

- We present CLiMRS, a multi-LLM cooperation framework for heterogeneous multi-robot collaboration that can perform long-horizon planning and efficient perception in complex tasks.
- We propose CLiMBench, a benchmark evaluating heterogeneous multi-robot collaboration with long-horizon assembly tasks, featuring varied skills and a realistic simulation environment.
- We demonstrate through extensive experiments that CLiMRS achieves significant efficiency improvements via dynamic group formation and cooperative long-horizon planning.

2 Related Works

2.1 Robotic Skills Training Across Scenarios

Single Agent Skill Training. Current approaches to train embodied skills for task execution generally follow two primary paradigms: rule-based and learning-driven methods. Traditional embodiment controllers optimize joint movements through the resolution of robotic kinematics, aiming to improve motion robustness and generate smoother, more precise trajectories (Kashyap and Parhi, 2021; Katayama et al., 2023). In recent years, advances in reinforcement learning and imitation learning have significantly improved robotic motion control across domains such as dexterous manipulation, bipedal locomotion, and quadrupedal navigation (Rajeswaran et al., 2017; Li et al., 2025; Bellegarda et al., 2024). These advances have progressively enabled embodied systems to coordinate actions in a cerebellum-like manner, supporting increasingly complex tasks in diverse environments.

Multi-agent Skill Training. Originally developed in game AI (Kurach et al., 2020; Perolat et al., 2022), multi-agent skill training has since been extended to more physically grounded domains such as robotics (Wang et al., 2024; Lai et al., 2025) and autonomous driving (Li et al., 2022). Despite these advances, multi-agent embodied learning remains limited by exponential state-space growth, and existing approaches such as mean-field methods (Yang et al., 2018) struggle to generalize across heterogeneous robots in collaboration.

To further this goal, we design a set of diverse robotic skills in CLiMBench that model robots’ low-level execution outcomes, capturing both successful executions and failure cases, thereby providing a testbed for planning methods in heterogeneous multi-agent collaboration.

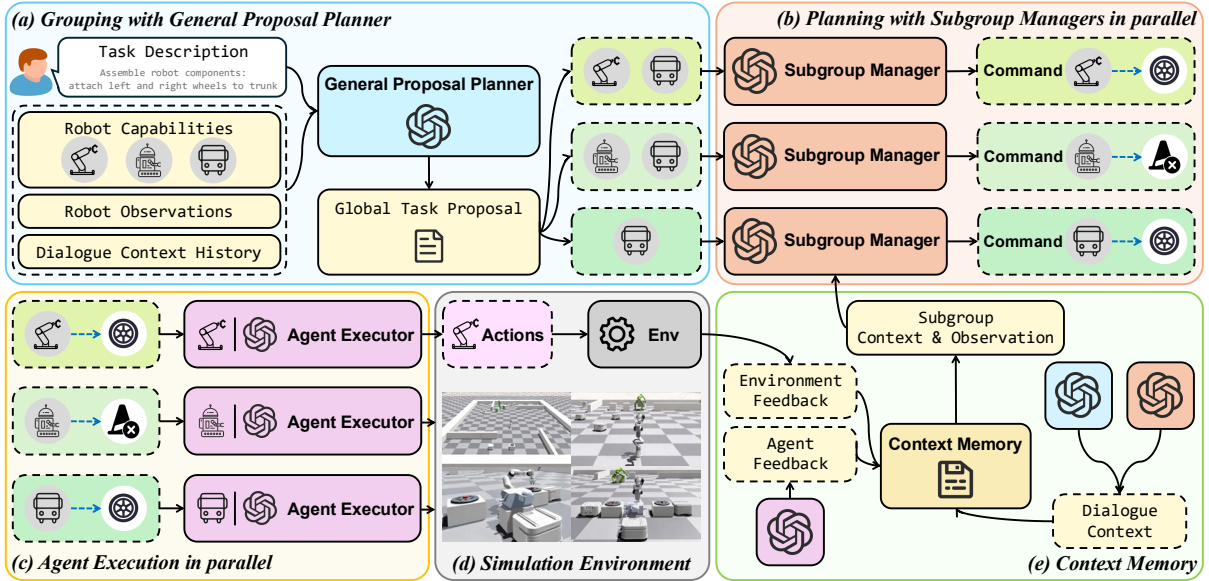


Figure 2: **CLiMRS Framework.** To employ our grouping–planning–execution–feedback cycle, CLiMRS comprises (a) a *general proposal planner* that forms dynamic agent subgroups, (b) multiple *subgroup managers* for local agent commands, (c) multiple *agent executors* for robot skills and agent feedback, (d) a *simulation environment* for environment interaction and feedback, and (e) a *context memory* module for all dialogue context and feedback.

2.2 Task Planning with LLMs in Robotics

LLM Planner for Robotics. The rapid progress of LLMs has motivated their use as task planners for robots, as few-shot and zero-shot learning enable the decomposition of high-level goals into executable task sequences (Brown et al., 2020; Huang et al., 2022; Kojima et al., 2022). Their code-generation abilities further allow LLMs to produce executable skill-level commands for robotic control (Liang et al., 2023a; Singh et al., 2023; Wang et al., 2023; Wu et al., 2023), while value-function-based methods guide robust skill selection (Lin et al., 2023; Ahn et al., 2022). Recent advances in prompting strategies further improve long-horizon LLM-based planning (Zhang et al., 2024b; Liu et al., 2025). Some studies also explore reasoning in joint or 3D action spaces (Mandi et al., 2023).

Multi-LLM Task Planning. Prior work emphasizes multi-LLM cooperation through discussion, debate, and role assignment (Chen et al., 2023a; Liang et al., 2023b; Hong et al., 2024) to provide complementary perspectives and improve output reliability, often augmented with feedback and memory mechanisms to improve long-horizon planning (Mandi et al., 2023; Liu et al., 2025; Zhang et al., 2024b; Wang et al., 2023).

Decision-making Paradigms in Multi-Robot Collaboration. Two primary paradigms emerged for complex multi-robot tasks: centralized and de-

centralized approaches. In decentralized schemes, multiple models or agents communicate, exchange intermediate plans, and iteratively refine their decisions through structured dialogue (Mandi et al., 2023; Zhang et al., 2024b; Liu et al., 2024b), while centralized methods typically rely on a single LLM to decompose global objectives and allocate tasks when planning (Kannan et al., 2023; Liu et al., 2025). A recent comparative study conducted in four diverse multi-agent scenarios (Chen et al., 2023b) further reports that centralized communication consistently achieves higher success rates and markedly greater token efficiency with proper feedback mechanisms, highlighting its strong potential for scalable real-world deployment.

Based on previous findings, CLiMRS structures multi-LLM cooperation through a grouping–planning–execution–feedback loop. By dynamically forming subgroups for subtasks, the framework enables parallel planning, agent execution, and robust collaboration.

3 Method

In this section, we present CLiMRS, an adaptive group negotiation framework among multiple LLMs that enables multi-robot collaboration. In human team problem-solving, a common and effective strategy for tackling complex tasks is to decompose the problem and assign subtasks to subgroups

that work in parallel. Inspired by this, our approach forms dynamic agent subgroups that hold centralized discussions on robot perception in parallel, with each robot paired with an independent LLM agent to give feedback to discussions, thus forming a dynamic grouping–planning–execution–feedback cycle, which is shown in Fig. 1.

As illustrated in Fig. 2, CLiMRS comprises five core modules: (a) a *general proposal planner* that forms dynamic agent groups, (b) multiple *subgroup managers* that generate local agent commands, (c) multiple *agent executors* that produce robot skills and return agent-level execution feedback, (d) a *simulation environment* for real-time interaction and provides environment feedback, and (e) a *context memory* module that records all inter-agent dialogues and feedback.

3.1 Grouping with General Proposal Planner

The first stage of our cycle dynamically partitions agents into subgroups responsible for different aspects of the overall task, using a *general proposal planner* to orchestrate the grouping process.

General Proposal Planner. As illustrated in Fig. 2(a), the *general proposal planner* generates a global task proposal that clusters agents into subtask-oriented teams. Given the overall task instruction, the prompted LLM incorporates robot capabilities, current observations, and dialogue history through a structured prompt. It outputs a plan with the following components: (1) *Situation Analysis*, assessing the environment and task progress; (2) *Spatial Analysis*, accounting for agent and object locations and spatial constraints; (3) *Task Decomposition*, breaking the objective into executable subtasks; (4) *Grouping Strategy*, clustering agents for parallel execution while minimizing interference; (5) *Subgoal Assignment*, specifying each group’s objective; (6) *Coordination Strategy*, outlining inter-group synchronization and execution order; and (7) *Risk Assessment*, identifying potential conflicts and mitigation plans. The resulting group-to-subtask assignments are then passed to the perception and execution modules.

3.2 Planning with Local Subgroup Managers

Given the agent groupings and their assigned subtasks, the second stage generates precise robot-level commands based on robot capabilities and current observations. As the subtasks are independent, multiple *subgroup managers* operate in parallel, as illustrated in Fig. 2(b).

Subgroup Manager. Each *subgroup manager* conducts a centralized discussion within its subgroup to determine fine-grained commands for individual robots. Conditioned on subtask instructions, subgroup robot capabilities, subgroup observations, and dialogue history, the subgroup manager selects appropriate skills and parameters for each agent.

3.3 Execution and Environment Interaction

With commands issued to robots, the final two stages of our cycle require them to evaluate these commands, determine appropriate actions, execute safely and provide feedback to refine future planning. The *agent executor* LLM verifies the feasibility of its command and issues the corresponding action only when the command is deemed executable.

Agent Execution. Shown in Fig. 2(c), the *agent executors* verify and execute commands from the *subgroup manager* while providing feedback. Each *agent executor* LLM considers its robot’s capabilities, current observations, and available actions. The executor first checks feasibility against the robot’s physical constraints and conditions. If feasible, the action is executed using the robot’s skills; otherwise, the robot remains idle for this cycle.

Simulation Environment. Shown in Fig. 2(d), the simulation environment serves as the execution backbone of our framework. It receives the robot skill execution signals issued by the *agent executors* and immediately carries out the corresponding low-level actions in real time. During execution, it monitors the evolving state of the environment and produces both updated robot observations and environment-level feedback. These outputs are fed back to the *context memory* shown in Fig. 2(e), allowing the overall method to track task progress, refine its environment understanding, and supply the information required for the next round of the grouping–planning–execution–feedback cycle.

3.4 Context Memory

With the grouping–planning–execution–feedback cycle described above, our framework employs a *context memory* module that contains execution feedback, environment observations, and dialogue history, providing context to all LLM agents.

Feedback Formation. Execution feedback is collected from two sources: (1) *environment feedback*, including updated environmental observations after robot actions are executed in the simulator, and (2) *agent feedback* generated by the *agent executors*. Agent executors report execution

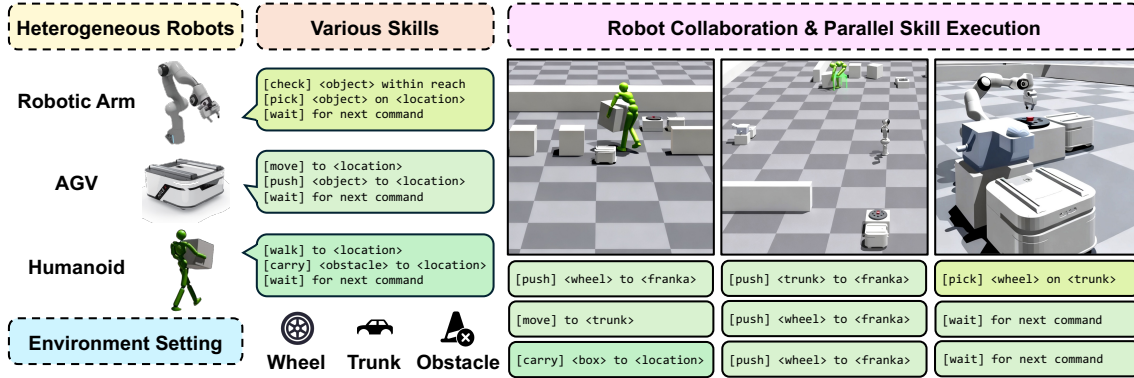


Figure 3: **Our Benchmark.** CLiMBench is a heterogeneous multi-robot collaboration benchmark, featuring multi-agent robots with diverse skills, enabling collaboration on complex assembly tasks of varying difficulty levels.

outcomes and provide diagnostic explanations for infeasible commands, enabling subsequent planners to reason about failures and task progress.

Context Memory Organization. As illustrated in Fig. 2(d), the *context memory* module stores: (1) *agent feedback*, including execution outcomes and diagnostic explanations from the *agent executors*; (2) *environment feedback*, consisting of updated observations from the *simulation environment* after action execution; and (3) *dialogue context*, including current planning dialogue and dialogue history accumulated across previous cycles.

The context memory provides tailored context inputs for different modules. For the *general proposal planner*, it retains the most recent five dialogue turns along with the latest environment observations, enabling agent feedback to inform new task proposals and subgroup formations. For the *subgroup managers*, it stores each subgroup’s latest observations and the last five dialogue turns, supplying rich, localized context to guide fine-grained planning and perception within each subgroup.

4 Benchmark

In this section, we introduce CLiMBench, a benchmark for heterogeneous multi-robot collaboration. We consider an assembly-oriented setting in which multiple components must be collected and assembled into a wheeled robot within a large workspace. Unlike prior multi-agent collaboration benchmarks (Liu et al., 2025), which decouple robot skill execution from the planning loop and assume successful execution by default, CLiMBench executes every robot skill within a realistic physics simulator, IsaacGym (Makoviychuk et al., 2021). This design explicitly allows execution failures and enables systematic evaluation of planning modules

under realistic execution uncertainty. The goal of CLiMBench is to evaluate long-horizon heterogeneous multi-robot collaboration in settings where planning and execution are tightly coupled.

4.1 Task Setting and Scene Construction

Task Setting. The goal of the task is to assemble the components into a wheeled robot. Within the workspace, the components are distributed across different rooms, some of which are occluded or blocked by obstacles that must be cleared before access is possible. As a result, robots must search for components, navigate complex environments, and perform coordinated manipulation under partial observability and execution uncertainty. This long-horizon task is considered successfully complete once all required components are correctly assembled into the target product within a step count limit to evaluate task execution efficiency.

Scene Construction. We instantiate this setting in the IsaacGym simulator by constructing the scene with multiple robotic agents and modular components, as is shown in Fig. 3. The environment includes three types of heterogeneous robots: one robotic arm, three autonomous ground vehicles (AGV), and one humanoid robot, each providing complementary capabilities for manipulation, transportation, and interaction with the environment to complete the assembly task. To decouple high-level planning from low-level control while keeping heterogeneous robot capabilities, robot actions are executed through predefined skills, allowing high-level planners to reason over realistic execution outcomes while abstracting low-level control.

Scene Initialization and Randomization. We initialize the environment and introduce controlled variations in task parameters and object configu-

Robot type	Skill list
Robotic Arm	[check] <i><object></i> within reach [pick] <i><object></i> on <i><location></i> [wait] for next command
AGV	[move] to <i><location></i> [push] <i><object></i> to <i><location></i> [wait] for next command
Humanoid	[walk] to <i><location></i> [carry] <i><obstacle></i> with both hands [wait] for next command

Table 1: **Robot Skills in CLiMBench.** We assign a distinct set of skills for each robot based on its specific capabilities in the assembly tasks in CLiMBench.

rations to enhance generalization. At the start of each episode, robots execute their skills under randomized task conditions, leading to diverse skill sequences and varying levels of agent-wise interaction. This provides a robust testbed for evaluating the effectiveness of different frameworks in multi-agent collaborative tasks.

Environment feedback. We design the environment feedback along two dimensions: (1) We update the state of all agents and the coordinates of objects within their perceptible range, and (2) We report conflicts that arise when multiple robots execute skills simultaneously.

4.2 Robot Skill Design in CLiMBench

In CLiMBench, each robot receives both the global task objectives and its observations (e.g., a humanoid robot observes its joint states, torso status, and target positions). As is shown in Table 1, we design distinct skills for different types of robots. Further details are provided in the Appendix.

Robotic Arm Manipulation. We use a Franka Panda arm to instantiate robotic manipulation skills in CLiMBench. To balance speed and precision, we adopt a two-stage strategy consisting of a fast coarse motion followed by fine adjustment phase. The arm is controlled using an operational-space controller (OSC) with gravity compensation, producing compliant spring-damper behavior in task space (Narang et al., 2022). Smooth operational-space waypoints are generated via interpolation to ensure reliable execution.

AGV Transportation. We use a TRACER Mini robot to instantiate AGV transportation skills in CLiMBench. Navigation paths are planned using a Rapidly-exploring Random Tree (RRT) planner

that accounts for environmental obstacles. The resulting paths are executed via differential-drive control, enabling smooth turns during navigation. For final delivery, the AGV follows straight-line motion to ensure reliable transportation and accurate placement at the assembly location.

Humanoid Carrying. We instantiate humanoid carrying as a goal-conditioned control skill trained with Adversarial Motion Priors (AMP) (Peng et al., 2021) following the single-object manipulation paradigm in COHOI (Gao et al., 2024). Given a target object and a goal location, the policy produces whole-body locomotion and manipulation to (1) approach and grasp the object, (2) carry it while maintaining balance and collision-robust gait, and (3) place it at the goal. Training combines task rewards for reaching, holding, and placing, and AMP style rewards to encourage natural, stable motions and rapid recovery, yielding reliable execution under contact and minor disturbances.

5 Experiments

In this section, we present a comprehensive evaluation of CLiMRS to address the following questions:

- (1) Is CLiMRS effective for simple daily-life multi-robot collaboration tasks?
- (2) Can CLiMRS perform well in challenging multi-robot assembly tasks?
- (3) Through ablation studies, how critical are the individual components of CLiMRS?

We evaluate CLiMRS in two distinct environments: CLiMBench and a simpler heterogeneous multi-robot collaboration benchmark from COHERENT (Liu et al., 2025). For LLM use, we use *gpt-4-0125-preview* to align with the setting in COHERENT. For quantitative analysis, we use task Success Rate (SR) and Average Step (AS) as evaluation metrics in this paper.

5.1 Evaluating CLiMRS on Simple Daily-Life Multi-Robot Collaboration

To answer Question (1), we evaluate CLiMRS on the COHERENT benchmark, a simpler heterogeneous multi-robot benchmark that includes diverse tasks across five real-world scenes, but involves at most three heterogeneous robots and assumes perfect skill execution. We adopt its evaluation metrics and use the reported results as our baseline.

Results shown in Table 2 and 3 suggest that CLiMRS succeeds on nearly all COHERENT tasks

Method	Mono-type Task		Dual-type Task		Trio-type Task		Average	
	SR	AS	SR	AS	SR	AS	SR	AS
DMRS-1D	0.700	10.6	0.467	18.0	0.667	20.7	0.600	17.2
DMRS-2D	0.500	11.5	0.267	19.9	0.400	24.5	0.375	19.6
CMRS	0.900	7.9	0.533	16.4	0.533	22.2	0.625	16.5
Primitive MCTS	0.000	14.0	0.000	21.5	0.000	26.9	0.000	21.7
LLM-MCTS	0.700	10.2	0.067	20.9	0.000	26.9	0.200	20.5
COHERENT	0.900	7.4	1.000	11.9	1.000	16.1	0.975	12.4
CLiMRS(Ours)	0.900	6.8	1.000	11.5	1.000	13.1	0.975	10.9
Ground Truth (GT)	–	6.5	–	10.3	–	12.9	–	10.3

Table 2: **Comparison Across Task Types in the COHERENT Benchmark.** CLiMRS outperforms all the baselines, achieving the largest gain on the most challenging trio-type tasks.

Method	S1		S2		S3		S4		S5		Average	
	SR	AS	SR	AS	SR	AS	SR	AS	SR	AS	SR	AS
DMRS-1D	0.500	17.4	0.625	15.8	0.625	18.3	0.750	15.1	0.500	19.3	0.600	17.2
DMRS-2D	0.500	18.9	0.500	18.3	0.375	20.6	0.250	18.9	0.250	21.1	0.375	19.6
CMRS	0.875	13.1	0.625	16.6	0.625	18.5	0.375	18.1	0.625	15.9	0.625	16.5
Primitive MCTS	0.000	21.5	0.000	21.8	0.000	22.5	0.000	20.5	0.000	22.0	0.000	21.7
LLM-MCTS	0.250	20.0	0.250	20.4	0.250	21.3	0.125	19.9	0.125	20.9	0.200	20.5
COHERENT	1.000	13.1	1.000	11.4	1.000	11.9	1.000	11.4	0.875	14.0	0.975	12.4
CLiMRS(Ours)	1.000	10.8	1.000	10.4	1.000	11.8	1.000	10.4	0.875	11.4	0.975	10.9
Ground Truth (GT)	–	10.3	–	10.4	–	10.8	–	9.8	–	10.5	–	10.3

Table 3: **Comparison Across Scenes in the COHERENT Benchmark.** CLiMRS outperforms all the baselines in every scene, demonstrating its superior performance.

and achieves higher efficiency with fewer steps. This trend holds across every scene, demonstrating our CLiMRS’ superior performance. Notably, in the most challenging trio-type tasks, which require all three robots to collaborate, CLiMRS delivers the largest gain, reducing the Average Step count by 18.6%, indicating that our approach offers stronger improvements on more complex tasks.

5.2 Evaluating CLiMRS on CLiMBench with Robot Assembly Tasks

To answer Question (2), we evaluate CLiMRS on CLiMBench. Our baselines include the following:

- DMRS-1D: a variant of CoELA (Zhang et al., 2024b), this decentralized framework lets robots determine their next step through dialogue, with the final decision summarized by the last robot.
- CMRS: a primitive centralized system (Huang et al., 2022) that uses a single decision-making LLM to output executable actions, where all information is stored in the prompt.
- COHERENT: an approximately centralized baseline that combines an oracle planning LLM

with a feedback LLM, where dialogue history is used to adjust future perception and planning.

For quantitative evaluation, we fix the scene parameters and select four representative scenarios to evaluate all methods. For each scenario, we manually derive minimal-step solutions as ground-truth references. We observe that scenarios with heavier object occlusion require more ground-truth steps, naturally forming two levels of difficulty.

Due to stochastic skill execution in CLiMBench, each task is executed five times, and we report the mean Success Rate (SR) and Average Steps (AS). To control evaluation cost and runtime, a task is considered successful only if it is completed within twice the corresponding ground-truth step count.

Results in Table 4 show that CLiMRS achieves a 100% success rate on CLiMBench, outperforming all baselines. In addition, CLiMRS reduces the Average Steps (AS) by 44.30% compared to the best baseline, indicating substantially higher planning efficiency in long-horizon collaboration.

Notably, most baselines terminate due to timeout and fail to complete tasks within the step

Method	Task 1 (Easy)		Task 2 (Easy)		Task 3 (Hard)		Task 4 (Hard)		Average	
	SR	AS	SR	AS	SR	AS	SR	AS	SR	AS
DMRS-1D	0.000	15.0	0.000	15.0	0.000	19.0	0.000	19.0	0.000	17.0
CMRS	0.000	15.0	0.000	15.0	0.000	19.0	0.000	19.0	0.000	17.0
COHERENT	1.000	13.6	0.800	13.6	0.400	18.2	0.600	17.8	0.700	15.8
CLiMRS (Ours)	1.000	8.2	1.000	8.4	1.000	9.4	1.000	9.2	1.000	8.8
Ground Truth (GT)	–	7.0	–	7.0	–	9.0	–	9.0	–	8.0

Table 4: **Comparison Across Tasks in CLiMBench.** CLiMRS outperforms all our baselines and reduces the Average Step (AS) by over 40%.

Method	Task 1 (Easy)		Task 2 (Easy)		Task 3 (Hard)		Task 4 (Hard)		Average	
	SR	AS	SR	AS	SR	AS	SR	AS	SR	AS
CLiMRS w/o history	0.000	15.0	0.000	15.0	0.000	19.0	0.000	19.0	0.000	17.0
CLiMRS w/o feedback	0.200	14.8	0.200	14.8	0.200	18.8	0.200	18.8	0.200	16.8
CLiMRS w/o grouping	0.600	14.0	0.800	13.2	0.600	17.2	0.600	17.4	0.650	15.5
CLiMRS (Ours)	1.000	8.2	1.000	8.4	1.000	9.4	1.000	9.2	1.000	8.8
Ground Truth (GT)	–	7.0	–	7.0	–	9.0	–	9.0	–	8.0

Table 5: **Ablation Studies.** Removing dialogue history, feedback information, or the grouping stage significantly reduces both Success Rate (SR) and Average Step (AS).

budget, and the performance degradation of baselines from the COHERENT benchmark (Table 2) to CLiMBench (Table 4) demonstrates that CLiMBench provides a significantly more challenging evaluation setting for robot collaboration.

Recover from Execution Failures. From our experiments, we observe three types of execution failures: (1) *improper grouping*, where no robot in the assigned group can complete the subtask; (2) *incorrect agent selection*, where a valid subtask is assigned to a robot without the required capabilities; and (3) *state inconsistency*, where missing information or unmet preconditions prevent proper execution. Nevertheless, the natural-language feedback provided by the executors captures these failure signals and, through the context memory, enables downstream perception and planning to be refined, ultimately leading to successful task completion.

5.3 Ablation Studies on CLiMRS

To answer Question (3), we conduct an ablation study to assess the contribution of each component in CLiMRS by: (i) removing the dialogue history, (ii) removing the feedback information, and (iii) removing the grouping stage from the grouping–planning–execution–feedback cycle. We evaluate all variants on the same tasks and metrics as in Section 5.2, with results summarized in Table 5.

The results show that both dialogue history and feedback information are critical to the framework: removing either component leads to failure in task completion, indicating that sustained context accumulation and execution-aware feedback are indispensable for long-horizon planning. The grouping stage also plays a central role in CLiMRS, significantly improving task success rates and reducing average steps by enabling effective coordination and parallelization across heterogeneous robots.

6 Conclusion

In this paper, we present CLiMRS, a human-team-inspired adaptive group negotiation framework among multiple LLMs for heterogeneous multi-robot collaboration, together with CLiMBench, a challenging benchmark for long-horizon assembly tasks under realistic execution uncertainty. Extensive experiments show that CLiMRS consistently outperforms all baselines in both success rate and efficiency, while robustly recovering from execution failures with feedback-driven negotiation. Ablation studies highlight the complementary roles of dialogue history, feedback, and dynamic subgroup formation. Overall, our results demonstrate that human-inspired group formation and negotiation principles offer an effective and scalable solution for heterogeneous multi-robot collaboration.

7 Limitations

In this paper, we primarily focus on improving the efficiency of multi-robot collaboration at the level of high-level behavioral planning. We assume that LLM inference latency is negligible compared to skill execution time, and therefore do not explicitly model network delays or computational costs in the current framework. Managing API costs under inference-efficiency constraints and exploring asynchronous inference–execution pipelines remain important directions for future work. In addition, while CLiMBench emphasizes long-horizon behavioral planning and coordination, it simplifies sensory inputs and perception modules in simulation. Bridging this gap toward more realistic perception and enabling sim-to-real transfer constitute promising avenues for future research.

8 Ethical Considerations

Our study investigates multi-robot collaboration using large language models (LLMs) for planning and negotiation in simulated environments. We discuss the primary ethical considerations relevant to this work below.

- **No Human or Sensitive Data.** Our research is conducted entirely in simulation and does not involve human subjects, personally identifiable information, or sensitive real-world data. We do not collect or process user data, nor do we rely on proprietary datasets.
- **Safety and Responsible Deployment.** While our methods and benchmark are evaluated exclusively in simulation, real-world deployment of autonomous multi-robot systems may pose physical safety risks. In particular, LLM-based planning may be susceptible to erroneous or inconsistent decisions under distribution shift. Any future deployment should therefore incorporate rigorous safety validation, fail-safe mechanisms, and human oversight, and comply with applicable safety standards and regulations.
- **Bias and Model Limitations.** The LLMs used in our study are pretrained by third parties and may reflect societal biases present in their training data. Although our work does not directly deploy LLMs in human-facing decision-making, such biases could indirectly influence planning behavior. We acknowledge this limitation and recommend bias auditing and robustness testing before any real-world application.

References

- Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, and 26 others. 2022. Do as i can and not as i say: Grounding language in robotic affordances. In *arXiv preprint arXiv:2204.01691*.
- Guillaume Bellegarda, Milad Shafiee, and Auke Ijspeert. 2024. Visual cpg-rl: Learning central pattern generators for visually-guided quadruped locomotion. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1420–1427. IEEE.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, and 12 others. 2020. [Language models are few-shot learners](#). *Preprint*, arXiv:2005.14165.
- Justin Chih-Yao Chen, Swarnadeep Saha, and Mohit Bansal. 2023a. Reconcile: Round-table conference improves reasoning via consensus among diverse llms. *arXiv preprint arXiv:2309.13007*.
- Yongchao Chen, Jacob Arkin, Yang Zhang, Nicholas Roy, and Chuchu Fan. 2023b. Scalable multi-robot collaboration with large language models: Centralized or decentralized systems? *arXiv preprint arXiv:2309.15943*.
- Jiawei Gao, Ziqin Wang, Zeqi Xiao, Jingbo Wang, Tai Wang, Jinkun Cao, Xiaolin Hu, Si Liu, Jifeng Dai, and Jiangmiao Pang. 2024. Coohei: Learning cooperative human-object interaction with manipulated object dynamics. In *Advances in Neural Information Processing Systems*.
- Sirui Hong, Mingchen Zhuge, Jonathan Chen, Xiawu Zheng, Yuheng Cheng, Jinlin Wang, Ceyao Zhang, Zili Wang, Steven Ka Shing Yau, Zijuan Lin, Liyang Zhou, Chenyu Ran, Lingfeng Xiao, Chenglin Wu, and Jürgen Schmidhuber. 2024. [MetaGPT: Meta programming for a multi-agent collaborative framework](#). In *The Twelfth International Conference on Learning Representations*.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. *arXiv preprint arXiv:2201.07207*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. Livecodebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint arXiv:2403.07974*.

676	Shyam Sundar Kannan, Vishnunandan LN Venkatesh, and Byung-Cheol Min. 2023. Smart-llm: Smart multi-agent robot task planning using large language models. <i>arXiv preprint arXiv:2309.10062</i> .	Kevin Lin, Christopher Agia, Toki Migimatsu, Marco Pavone, and Jeannette Bohg. 2023. Text2motion: from natural language instructions to feasible plans . <i>Autonomous Robots</i> .	732 733 734 735
680	Abhishek Kumar Kashyap and Dayal R Parhi. 2021. Particle swarm optimization aided pid gait controller design for a humanoid robot. <i>ISA transactions</i> , 114:306–330.	Kehui Liu, Zixin Tang, Dong Wang, Zhigang Wang, Xuelong Li, and Bin Zhao. 2025. Coherent: Collaboration of heterogeneous multi-robot system with large language models. In <i>2025 IEEE International Conference on Robotics and Automation (ICRA)</i> , pages 10208–10214. IEEE.	736 737 738 739 740 741
684	Sotaro Katayama, Masaki Murooka, and Yuichi Tazaki. 2023. Model predictive control of legged and humanoid robots: models and algorithms. <i>Advanced Robotics</i> , 37(5):298–315.	Xinzhu Liu, Peiyan Li, Wenju Yang, Di Guo, and Huaping Liu. 2024a. Leveraging large language model for heterogeneous ad hoc teamwork collaboration. <i>arXiv preprint arXiv:2406.12224</i> .	742 743 744 745
688	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In <i>Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22</i> , Red Hook, NY, USA. Curran Associates Inc.	Xinzhu Liu, Peiyan Li, Wenju Yang, Di Guo, and Huaping Liu. 2024b. Leveraging large language model for heterogeneous ad hoc teamwork collaboration . <i>Preprint</i> , arXiv:2406.12224.	746 747 748 749
694	Karol Kurach, Anton Raichuk, Piotr Stańczyk, Michał Zajac, Olivier Bachem, Lasse Espeholt, Carlos Riquelme, Damien Vincent, Marcin Michalski, Olivier Bousquet, and 1 others. 2020. Google research football: A novel reinforcement learning environment. In <i>Proceedings of the AAAI conference on artificial intelligence</i> , volume 34, pages 4501–4510.	Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. 2019. Amass: Archive of motion capture as surface shapes. In <i>Proceedings of the IEEE/CVF international conference on computer vision</i> , pages 5442–5451.	750 751 752 753 754
701	Matthew Lai, Keegan Go, Zhibin Li, Torsten Kröger, Stefan Schaal, Kelsey Allen, and Jonathan Scholz. 2025. Roboballet: Planning for multirobot reaching with graph neural networks and reinforcement learning. <i>Science Robotics</i> , 10(106):eads1204.	Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and 1 others. 2021. Isaac gym: High performance gpu-based physics simulation for robot learning. <i>arXiv preprint arXiv:2108.10470</i> .	755 756 757 758 759 760
706	Yiming Li, Dekun Ma, Ziyang An, Zixun Wang, Yiqi Zhong, Siheng Chen, and Chen Feng. 2022. V2x-sim: Multi-agent collaborative perception dataset and benchmark for autonomous driving. <i>IEEE Robotics and Automation Letters</i> , 7(4):10914–10921.	Zhao Mandi, Shreeya Jain, and Shuran Song. 2023. Roco: Dialectic multi-robot collaboration with large language models . <i>Preprint</i> , arXiv:2307.04738.	761 762 763
711	Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. 2025. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. <i>The International Journal of Robotics Research</i> , 44(5):840–888.	Zhao Mandi, Shreeya Jain, and Shuran Song. 2024. Roco: Dialectic multi-robot collaboration with large language models. In <i>2024 IEEE International Conference on Robotics and Automation (ICRA)</i> , pages 286–299. IEEE.	764 765 766 767 768
716	Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. 2023a. Code as policies: Language model programs for embodied control . In <i>2023 IEEE International Conference on Robotics and Automation (ICRA)</i> , pages 9493–9500.	Christopher E Mower, Yuhui Wan, Hongzhan Yu, Antoine Grosnit, Jonas Gonzalez-Billandon, Matthieu Zimmer, Jinlong Wang, Xinyu Zhang, Yao Zhao, Anbang Zhai, and 1 others. 2024. Ros-llm: A ros framework for embodied ai with task feedback and structured reasoning. <i>arXiv preprint arXiv:2406.19741</i> .	769 770 771 772 773 774
722	Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Zhaopeng Tu, and Shuming Shi. 2023b. Encouraging divergent thinking in large language models through multi-agent debate. <i>arXiv preprint arXiv:2305.19118</i> .	Yashraj Narang, Kier Storey, Ireteyayo Akinola, Miles Macklin, Philipp Reist, Lukasz Wawrzyniak, Yunrong Guo, Adam Moravanszky, Gavriel State, Michelle Lu, Ankur Handa, and Dieter Fox. 2022. Factory: Fast contact for robotic assembly . <i>Preprint</i> , arXiv:2205.03532.	775 776 777 778 779 780
727	Wenlong Liang, Rui Zhou, Yang Ma, Bing Zhang, Songlin Li, Yijia Liao, and Ping Kuang. 2025. Large model empowered embodied ai: A survey on decision-making and embodied learning. <i>arXiv preprint arXiv:2508.10399</i> .	Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. 2021. Amp: adversarial motion priors for stylized physics-based character control . <i>ACM Transactions on Graphics</i> , 40(4):1–20.	781 782 783 784

785	Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T Connor, Neil Burch, Thomas Anthony, and 1 others. 2022. Mastering the game of stratego with model-free multiagent reinforcement learning. <i>Science</i> , 378(6623):990–996.	840
786		841
787		842
788		843
789		844
790		
791	Aske Plaat, Annie Wong, Suzan Verberne, Joost Broekens, Niki van Stein, and Thomas Bäck. 2024. Reasoning with large language models, a survey. <i>CoRR</i> .	845
792		846
793		847
794		848
795	Aravind Rajeswaran, Vikash Kumar, Abhishek Gupta, Giulia Vezzani, John Schulman, Emanuel Todorov, and Sergey Levine. 2017. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. <i>arXiv preprint arXiv:1709.10087</i> .	849
796		850
797		851
798		852
799		
800		
801	David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In <i>First Conference on Language Modeling</i> .	853
802		854
803		855
804		856
805		857
806	Sahar Salimpour, Lei Fu, Farhad Keramat, Leonardo Militano, Giovanni Toffetti, Harry Edelman, and Jorge Peña Queralta. 2025. Towards embodied agentic ai: Review and classification of llm-and vlm-driven robot autonomy and interaction. <i>arXiv preprint arXiv:2508.05294</i> .	
807		
808		
809		
810		
811		
812	John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. <i>arXiv preprint arXiv:1707.06347</i> .	
813		
814		
815		
816	Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, and Animesh Garg. 2023. Prog-prompt: Generating situated robot task plans using large language models . In <i>2023 IEEE International Conference on Robotics and Automation (ICRA)</i> , pages 11523–11530.	
817		
818		
819		
820		
821		
822		
823	Chan Hee Song, Jiaman Wu, Clayton Washington, Brian M Sadler, Wei-Lun Chao, and Yu Su. 2023. Llm-planner: Few-shot grounded planning for embodied agents with large language models. In <i>Proceedings of the IEEE/CVF international conference on computer vision</i> , pages 2998–3009.	
824		
825		
826		
827		
828		
829	Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. 2023. Voyager: An open-ended embodied agent with large language models. <i>arXiv preprint arXiv: Arxiv-2305.16291</i> .	
830		
831		
832		
833		
834	Weizheng Wang, Le Mao, Ruiqi Wang, and Byung-Cheol Min. 2024. Multi-robot cooperative socially-aware navigation using multi-agent reinforcement learning. In <i>2024 IEEE International Conference on Robotics and Automation (ICRA)</i> , pages 12353–12360. IEEE.	
835		
836		
837		
838		
839		
	Jimmy Wu, Rika Antonova, Adam Kan, Marion Lepert, Andy Zeng, Shuran Song, Jeannette Bohg, Szymon Rusinkiewicz, and Thomas Funkhouser. 2023. Tidybot: Personalized robot assistance with large language models. <i>Autonomous Robots</i> .	
	Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. 2018. Mean field multi-agent reinforcement learning. In <i>International conference on machine learning</i> , pages 5571–5580. PMLR.	
	Hangtao Zhang, Chenyu Zhu, Xianlong Wang, Ziqi Zhou, Shengshan Hu, and Leo Yu Zhang. 2024a. Badrobot: Jailbreaking llm-based embodied ai in the physical world. <i>arXiv preprint arXiv:2407.20242</i> , 3.	
	Hongxin Zhang, Weihua Du, Jiaming Shan, Qinhong Zhou, Yilun Du, Joshua B. Tenenbaum, Tianmin Shu, and Chuang Gan. 2024b. Building cooperative embodied agents modularly with large language models . <i>Preprint</i> , arXiv:2307.02485.	

Appendix

A Additional Details on CLiMBench

A.1 Simulation Setup and Agent Parameters

Physics Simulation. All tasks in CLiMBench are developed in the IsaacGym physics simulator (Makoviychuk et al., 2021). The simulator executes robot actions in a step-by-step manner, advancing the physical state of the environment after each issued skill. Gravity, collision detection, and joint limit constraints are enabled for all agents to ensure physically consistent execution.

Agent Roster and Physical Specifications. The assembly scene in CLiMBench includes three types of heterogeneous robotic agents: a single robotic manipulator, three autonomous ground vehicles (AGVs), and one humanoid. Table 6 summarizes the key specifications of each agent type, including degrees of freedom, payload capacity, maximum reach, and the number of instantiated agents.

	DoF	Payload	Max Reach	Quantity
Franka	7	3kg	855mm	1
TRACER Mini	3	80kg	All	3
Humanoid	28	–	All	1

Table 6: Agent Parameters in CLiMBench.

The robotic arm is instantiated as a Franka Panda manipulator with seven degrees of freedom and a maximum reach of 855 mm, suitable for precise manipulation and assembly. The AGVs are instantiated as TRACER Mini mobile robots with differential-drive kinematics and high payload capacity, enabling reliable transportation of components within the workspace. The humanoid features a high-dimensional articulated body with 28 degrees of freedom, supporting whole-body locomotion and object carrying skills.

A.2 Environment Setup and Randomization

Environment Workspace Initialization. All tasks in CLiMBench are executed within a shared global coordinate frame. The workspace domain, component placement candidates, and agent initialization ranges are summarized in Table 7. All object poses and agent base positions are specified in this global frame, while the vertical axis is used only to model object height and collision geometry.

Obstacle Configuration. In addition to randomly initialized agents and components, the environment includes static obstacles that block navigation and cause occlusions.

Entity	Initialization Range
Environment Domain	$(x, y) \in [-6, 6] \times [-10, 10]$
Components (Wheels & Trunk)	$(x, y) \in \{(4, 8), (-4, 8), (-4, -8), (4, -8)\}$
AGVs	$x \in [-3, 3], y \in [-5, 5]$
Humanoid	Central region of the environment

Table 7: Random Initialization of the Environment Components and Agents in CLiMBench.

Four rectangular obstacles of size 1×5 m are placed near the corners of the workspace. Three smaller obstacles of size $1 \times 1 \times 0.5$ m are positioned along primary traversal paths of the mobile agents, introducing potential collisions and blocked access regions.

Randomization Policy. For each episode, component locations and agent initial positions are independently sampled from the ranges specified in Table 7. All methods are evaluated under identical environment configurations and random seeds to ensure fair comparison. No curriculum learning or adaptive environment scheduling is employed.

A.3 Low-level Skill Implementation Details

Franka Manipulation Implementation. Franka manipulation skills are executed using a numerical inverse kinematics (IK) solver combined with an operational-space controller (OSC). The IK solver computes target joint configurations for the first seven joints, with a maximum of 200 iterations and a residual tolerance of 10^{-4} . The Franka base is initialized at a fixed pose $(0, -2.0, 0)$ with orientation quaternion $(0, 0, 0, 1)$. Task-space motion is controlled using an impedance-based OSC with:

$$\mathbf{F} = k_p \mathbf{e} + k_v \dot{\mathbf{e}}, \quad (1)$$

where the proportional gain is set to $k_p = 5$ and the derivative gain is chosen as $k_v = 2\sqrt{k_p}$ to achieve critical damping. The gripper closing threshold is set to 0.05 m, accounting for the object grasping radius of 0.03 m used in CLiMBench.

During assembly, components are aligned and attached using a magnetic attraction model with an effective range of 0.03 m. For each manipulation trajectory, operational-space waypoints are generated via interpolation to ensure smooth and continuous execution.

AGV Navigation and Motion Control. AGV transportation skills are implemented using a Rapidly-exploring Random Tree (RRT) planner operating in the planar workspace. The planner considers all static and dynamic obstacles and robots in the environment as the collision space. The

RRT planner uses a collision-checking resolution of 0.1 m along each edge and a rewire count of 32. Once a feasible path is found, it is executed using differential-drive control to enable smooth turning and stable navigation. For final positioning during delivery, the AGV follows a straight-line motion segment to improve placement reliability.

Humanoid Observation and Skill Interface.

Humanoid skills are executed by a pretrained policy that operates on a structured observation interface. Let $N_j = 15$ denote the number of articulated bodies of the humanoid. For each body $j \in \{1, \dots, N_j\}$, the observation is defined as

$$\mathbf{o}_j = [\mathbf{q}_j, \mathbf{p}_j, \mathbf{v}_j, \mathbf{a}_j], \quad (2)$$

where $\mathbf{q}_j \in \mathbb{R}^6$ represents the body relative rotational and translational configuration, $\mathbf{p}_j \in \mathbb{R}^3$ the Cartesian position, $\mathbf{v}_j \in \mathbb{R}^3$ the linear velocity, and $\mathbf{a}_j \in \mathbb{R}^3$ the linear acceleration. The full humanoid observation vector is constructed by concatenating all body-level observations and a task observation:

$$\mathbf{o} = [\mathbf{o}_1, \mathbf{o}_2, \dots, \mathbf{o}_{N_j}, \mathbf{o}_{task}]. \quad (3)$$

Observations of the root body and global reference frame are excluded, as they encode absolute pose information rather than local kinematic and dynamic states. This formulation provides a consistent interface for skill execution while remaining independent of the humanoid’s global position and orientation in the environment.

A.4 Humanoid Carry Skill Training Details

AMP Objective and Reward Design. Humanoid carrying skills are trained using the Adversarial Motion Priors (AMP) framework (Peng et al., 2021), following the single-object manipulation paradigm introduced in COOHOI (Gao et al., 2024). AMP augments the task objective with a discriminator-based style reward that encourages motions consistent with human motion data. At each time step t , the total reward is defined as

$$r_t = w^G r^G(\mathbf{s}_t, \mathbf{g}_t, \mathbf{s}_{t+1}) + w^S r^S(\mathbf{s}_t, \mathbf{s}_{t+1}), \quad (4)$$

where r^G denotes the task-specific reward and r^S is the AMP style reward. The task reward consists of components for walking, holding, and placement:

$$r^G = 0.2 r_{\text{walk}}^G + 0.4 r_{\text{held}}^G + 0.4 r_{\text{target}}^G, \quad (5)$$

while the style reward is computed as

$$r^S(\mathbf{s}_t, \mathbf{s}_{t+1}) = -\log(1 - D(\mathbf{s}_t, \mathbf{s}_{t+1})), \quad (6)$$

Parameters	Value
num envs	16384
episode length	600
discount factor	0.99
e_clip	0.2
horizon_length	32
minibatch_size	16384
amp_minibatch_size	4096
task_reward_weight	0.5
disc_reward_weight	0.5

Table 8: Humanoid Carrying Skill Training.

with $D(\cdot)$ denoting the discriminator trained on reference human motion data.

Discriminator Training. The AMP discriminator is trained to distinguish between policy-generated motion transitions and reference human motion data. We adopt the same discriminator architecture and training procedure as in (Peng et al., 2021), where the discriminator operates on pairs of consecutive states $(\mathbf{s}_t, \mathbf{s}_{t+1})$ and outputs the probability that the transition originates from real motion data. Reference motions are drawn from the ACCAD subset of the AMASS dataset (Mahmood et al., 2019), which contains a diverse set of human motions including locomotion, lifting, carrying, and object placement. During training, the discriminator and the policy are updated alternately. The discriminator is trained using binary cross-entropy loss to separate reference transitions from policy-generated transitions, while the policy receives the discriminator-based style reward defined in Eq. 6. This adversarial training scheme encourages the humanoid policy to produce stable and human-like motions while optimizing task-specific objectives.

Domain Randomization. To improve robustness and generalization in assembly environments, we apply domain randomization during training. Randomized factors include object orientation, relative alignment between the humanoid and the object, and object proximity. These variations encourage the learned policy to tolerate spatial uncertainty and contact perturbations commonly encountered during execution in CLiMBench.

Training Hyperparameters. The humanoid policy is trained using Proximal Policy Optimization (PPO) (Schulman et al., 2017). Key training hyperparameters, including the number of parallel environments, episode length, reward weights, and optimization settings, are summarized in Table 8.

Training Infrastructure. Training is conducted

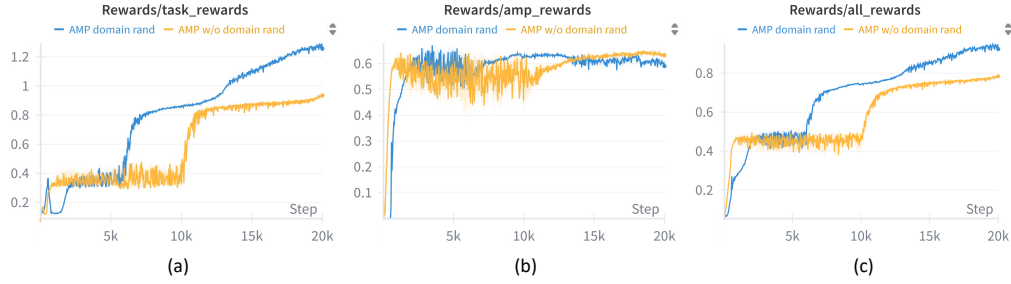


Figure 4: **Ablation Study on AMP Domain Randomization.** Results show that Domain Randomization helps the humanoid agent to obtain a stable and reusable carrying skill in CLiMBench.

Single Agent Prompt Templates

Think of yourself as one of the robots `#ID(101, 606, 202, 203, or 204)`, and your available skills are:

Available Skills for Wheeled Robots:
 - `#[SKILL#] <#ROBOT#> (#ID) move to component location using RRT path`

Your role: You are a `#character skill description#`

The list of available actions you can perform in your current environment is: `#ACTIONLIST#`

The instructions for the next step of the task is: `#INSTRUCTION#`

Please choose the best available action to achieve the goal as soon as possible.

Note: If there is an action in the available actionlist that can satisfy the instruction, the first sentence in the output needs to be "YES I CAN."

Note: if there is not an action in the available actionlist that can satisfy the instruction, then *output the possible reasons*, such as the agent does not have the ability to execute the current instruction and needs the assistance of other types of agents; or the current agent is not in the state of executing this instruction and an additional action needs to be performed as a prerequisite before the instruction can be executed; or the current state already meets the requirements of the task instructions and no other actions need to be performed. The first sentence in the output needs to be "SORRY I CANNOT."

Note: Think of yourself as `#ROBOT#`, generating content in a first-person conversation.

Let's think step by step, but you should strictly obey the rules above about the first sentence in the output, do not include "*" in the first sentence to affect output parsing.

Figure 5: **Single Agent Prompt Templates.** This template encodes the agent's name, unique identifier, and action specification to ensure structured and unambiguous command interpretation for an *agent executor*.

1026 in IsaacGym headless mode on a single NVIDIA
 1027 A100 GPU to enable large-scale parallel simulation.
 1028 After training, the learned policy is deployed with
 1029 rendering enabled on a workstation equipped with
 1030 an NVIDIA RTX 4060 Ti for evaluation.

1031 **Ablation on Domain Randomization.** Fig. 4
 1032 presents an ablation on domain randomization in
 1033 AMP training. Without randomization, it con-
 1034 verges to a local optimum early and exhibits poor
 1035 generalization, while domain-randomized training
 1036 achieves consistently higher task rewards and more
 1037 reliable execution. This result supports the use
 1038 of domain randomization to obtain a stable and
 1039 reusable carrying skill for the benchmark.

1040 A.5 LLM-to-Skill Execution Interface

1041 **Command Schema and Prompt Templates.** All
 1042 robot actions in CLiMBench are issued through
 1043 structured natural-language commands generated

1044 by LLM-based planners. Each command explicitly
 1045 specifies the action type, the acting agent(s), the
 1046 target object, and the target location. Figure 5 il-
 1047 lustrates the single-agent prompt templates used to
 1048 ensure parsable command generation. These tem-
 1049 plates encode the agent name, unique identifier, and
 1050 action specification, providing a unified interface
 1051 across heterogeneous robots.

1052 **Command Parsing via Regular Expressions.**
 1053 Given an LLM-generated command C , the execu-
 1054 tion interface parses it into a structured representa-
 1055 tion using a set of predefined regular expressions.
 1056 Formally, each command is mapped as

$$1057 C \xrightarrow{\text{parse}} \{G, A, O, L\}, \quad (7)$$

1058 where G denotes the group identifier, $A =$
 1059 $\{(a_i, id_i)\}$ is the set of assigned agents with their
 1060 identifiers, O specifies the target object and its iden-
 1061 tifier, and L denotes the target location. Only com-

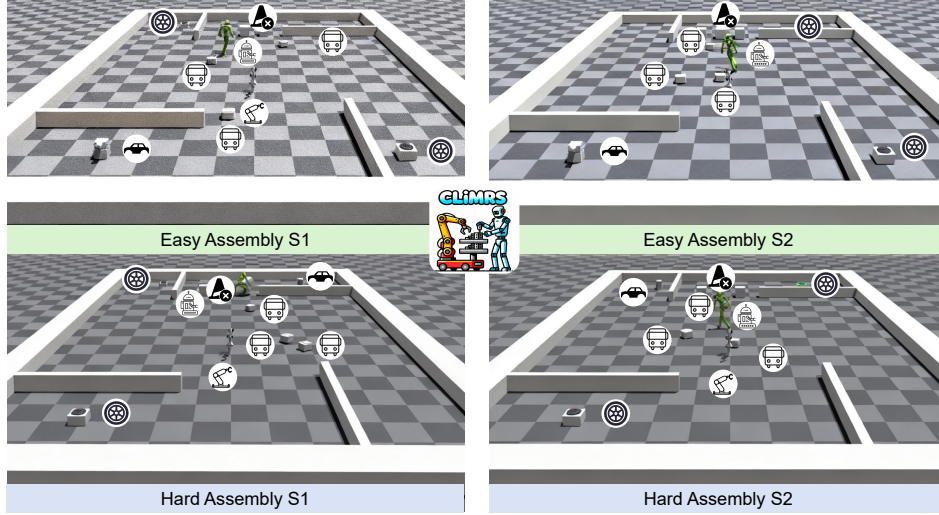


Figure 6: **Tasks of Different Difficulties.** To evaluate the effectiveness of our method in facilitating multi-robot collaboration, we design four heterogeneous agent scenarios under two difficulty tiers, namely *easy* and *hard*.

mands that can be successfully parsed into this structure are considered for further verification.

Command Verification. To ensure reliable and safe execution, each parsed command is validated through a four-stage verification pipeline. Let f denote an LLM function call corresponding to a structured command. The verification operator V maps f to a validation vector

$$f \xrightarrow{V} \{v_1, v_2, v_3, v_4\}, \quad (8)$$

where each component corresponds to a specific verification stage. The pipeline consists of the following four stages:

- **Capability assessment.** The system first evaluates whether the assigned agent(s) are capable of executing the requested action based on their available skills. The assessment yields a binary decision or confidence score, and commands below a preset threshold τ_c are rejected or deferred.
- **Action selection.** For commands that pass the capability assessment, the planner selects a concrete action a^* from the available action set \mathcal{A} , optionally providing a ranked subset or per-action confidence.
- **Action parsing.** The selected action a^* is processed by a parsing operator \mathcal{P} to produce a structured command $C^{\text{struct}} = \mathcal{P}(a^*)$, which includes fields such as action type, agent assignment, target object, and target location.
- **Execution verification (judge).** A dedicated judge module \mathcal{J} evaluates C^{struct} for format

correctness, semantic consistency, physical feasibility, and execution safety.

Only commands that satisfy all verification criteria are forwarded to the execution module. Overall, the verification operator can be expressed as

$$V = \mathcal{J} \circ \mathcal{P} \circ \mathcal{S}_2 \circ \mathcal{S}_1, \quad (9)$$

where \mathcal{S}_1 and \mathcal{S}_2 denote the capability assessment and action selection stages, respectively.

Failure Handling and Feedback Logging. Commands that fail at any stage of the verification pipeline are not executed. Instead, the interface records structured failure feedback indicating the reason for rejection, such as insufficient agent capability, semantic mismatch, or physical infeasibility. This feedback is propagated to the future planning process and stored in the *context memory* module, enabling planners to adjust subsequent decisions. Deferred commands may be reconsidered in later planning iterations once agent states change or environment conditions are ready.

A.6 Benchmark Evaluation Details

Task Difficulty Settings. We define two levels of task difficulty, referred to as *easy* and *hard*, which differ in obstacle configuration and agent initialization. In the easy setting, AGV transportation paths are free of blocking obstacles, and initial agent placements are arranged to avoid potential collisions with the humanoid. In contrast, the hard setting introduces obstacles along AGV trajectories

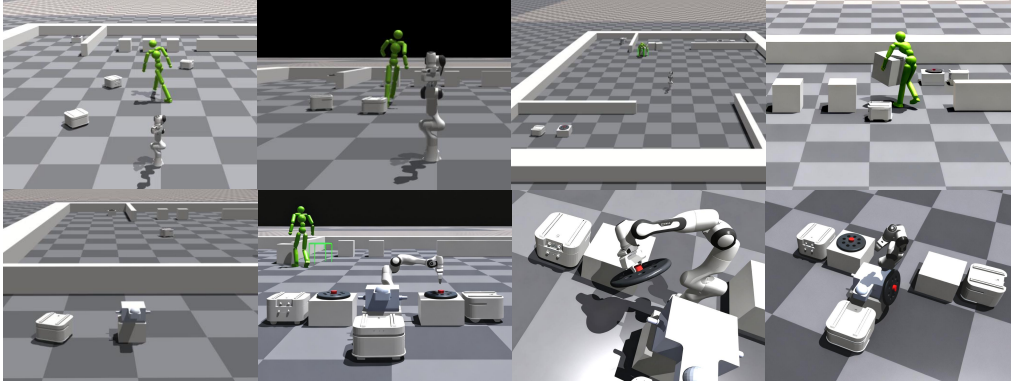


Figure 7: **Simulation.** It shows an example of the assembly process in the dynamic simulation.

Parallel Execution

```

Group 0: <humanoid>(101) - Sub-goal: Clear obstacles to ensure paths are clear for component transportation
Group 1: <mobile_car_1>(201) - Sub-goal: Transport the left wheel (405) to the assembly area
Group 2: <mobile_car_2>(202) - Sub-goal: Transport the right wheel (406) to the assembly area
Group 3: <mobile_car_3>(203) - Sub-goal: Transport the trunk (303) to the assembly area
Group 4: <robot arm>(606) - Sub-goal: Assemble the components by attaching the wheels to the trunk once all components
are delivered
Processing group 0 message: Hello <humanoid>(101): Please use [carry] skill to clear obstacles.
#####
the first sentence is YES I CAN
#####
No more things to do!
Processing group 1 message: Hello <mobile_car_1>(201): Please use [move] skill to navigate to <left wheel>(405)
location, then use [push] skill to transport it to the franka assembly area.
#####
the first sentence is YES I CAN
#####
[move] <mobile_car_1> (201) move to assigned component location using RRT path
Processing group 4 message: "Hello <franka>(606): Please use [check] skill to check <trunk>(303) for assembly
readiness."
#####
the first sentence is YES I CAN
#####
[check] <franka> (606) check <trunk> (303)

```

Figure 8: **Parallel Execution.** To enhance task execution efficiency, agents within each group are executed in parallel, as reflected in the execution log.

1121 and initializes agents in closer proximity, increas- 1139
1122 ing the likelihood of collisions and requiring more 1140
1123 coordination during execution. Examples of the 1141
1124 two settings are illustrated in Fig. 6. 1142

1125 **Trials and Stochasticity.** Due to stochasticity in 1143
1126 low-level skill execution and environment interac- 1144
1127 tions, each task configuration is evaluated over mul- 1145
1128 tiple independent trials. Unless otherwise speci- 1146
1129 fied, we run each task five times and report aver- 1147
1130 aged metrics across trials. All methods are evalu- 1148
1131 ated under identical environment configurations and 1149
1132 random seeds to ensure fair comparison. An exam- 1150
1133 ple of the assembly process can be seen in Fig. 7. 1151

1134 **Step Budget and Timeout Criteria.** To control 1152
1135 evaluation cost and ensure comparability, each 1153
1136 task is assigned a maximum step budget. The step 1154
1137 budget is defined as twice the number of steps in 1155
1138 a manually derived minimal-step solution for the

1139 corresponding task. A task is considered unsuccess- 1140
1141 ful if the target assembly is not completed within 1141
1142 this budget, even if individual skills execute suc- 1142
1143 cessfully. This criterion penalizes inefficient long- 1143
1144 horizon planning and excessive replanning. 1144

1145 **Execution Logging.** During evaluation, the 1145
1146 benchmark records detailed execution logs, includ- 1146
1147 ing agent group assignments, issued skill com- 1147
1148 mands, verification outcomes, and environment 1148
1149 feedback at each step. These logs enable post-hoc 1149
1150 analysis of planning behavior, grouping dynamics, 1150
1151 and failure recovery, and are used for qualitative 1151
1152 visualization and debugging. An example of the 1152
1153 parallel execution logs can be seen in Fig. 8. 1153

B Example Logs and Prompts 1153

1154 We show some examples of logs and prompts in all 1154
1155 the figures below. 1155

Scene Dynamics Log

You are a structured output formatter for multi-agent task coordination. Your job is to convert a detailed grouping strategy into a precise, parseable format.

Step:0

Assemble robot components: attach left and right wheels to trunk

Group 0: <humanoid>(101) - Sub-goal: **Clear** obstacles to ensure paths are clear for component transportation

Group 1: <mobile_car_1>(201) - Sub-goal: Transport the left wheel (405) to the assembly area

Group 2: <mobile_car_2>(202) - Sub-goal: Transport the right wheel (406) to the assembly area

Group 3: <mobile_car_3>(203) - Sub-goal: Transport the trunk (303) to the assembly area

Group 4: <robot arm>(606) - Sub-goal: Assemble the components by attaching the wheels to the trunk once all components are delivered

Step:1

Assemble robot components: attach left and right wheels to trunk

Group 0: <humanoid>(101) - Sub-goal: **Clear** any obstacles that might prevent the **mobile cars** from reaching and transporting their components

Group 1: <mobile_car_1>(201) - Sub-goal: Transport the left wheel to the franka assembly area

Group 2: <mobile_car_2>(202) - Sub-goal: Transport the right wheel to the franka assembly area

Group 3: <mobile_car_3>(203) - Sub-goal: Transport the trunk to the franka assembly area

Group 4: <robot arm>(606) - Sub-goal: Assemble the components once delivered to the assembly area

Step:2

Assemble robot components: attach left and right wheels to trunk

Group 0: <humanoid>(101) - Sub-goal: Ensure all paths are clear for mobile cars to transport components

Group 1: <mobile_car_1>(201) - Sub-goal: Transport left wheel to the assembly area

Group 2: <mobile_car_2>(202) - Sub-goal: Transport right wheel to the assembly area

Group 3: <mobile_car_3>(203) - Sub-goal: Transport trunk to the assembly area

Group 4: <robot arm>(606) - Sub-goal: Assemble left and right wheels onto trunk when components are ready

Step:3

Assemble robot components: attach left and right wheels to trunk

Group 0: <mobile_car_1>(201), <mobile_car_2>(202), <mobile_car_3>(203) - Sub-goal: Reposition components if not in optimal position for assembly

Group 1: <robot arm>(606) - Sub-goal: **Check** readiness of left and right wheels (405 and 406) for assembly, then assemble onto trunk (303)

Group 2: <humanoid>(101) - Sub-goal: Clear any obstacles if necessary to ensure accessibility of assembly area

Step:4

Assemble robot components: attach left and right wheels to trunk

Group 0: <robot arm>(606) - Sub-goal: **Perform final check** on right wheel readiness and assemble components when ready

Group 1: <mobile_car_2>(202) - Sub-goal: **Push right wheel** to within franka's operational range for readiness check

Group 2: <mobile_car_1>(201), <mobile_car_3>(203) - Sub-goal: **Position** left wheel and trunk **within franka's** operational range for assembly

Group 3: <humanoid>(101) - Sub-goal: Remain on standby to clear any unforeseen obstacles

Step:5

Assemble robot components: attach left and right wheels to trunk

Group 0: <mobile_car_1>(201), <mobile_car_2>(202), <mobile_car_3>(203) - Sub-goal: Transport their designated components to the franka assembly area sequentially to avoid traffic and ensure all components are ready for assembly

Group 1: <robot arm>(606) - Sub-goal: Assemble the robot by attaching the left and right wheels to the trunk once all components are within its operational range

Group 2: <humanoid>(101) - Sub-goal: **Remain on standby** to clear any unforeseen obstacles that might arise during the transport of components

Step:6

Assemble robot components: attach left and right wheels to trunk

Group 0: <robot arm>(606) - Sub-goal: **Assemble** the robot by attaching the left and right wheels to the trunk

Group 1: <humanoid>(101) - Sub-goal: Standby for path clearing and obstacle management if unforeseen obstacles appear

Step:7

Assemble robot components: attach left and right wheels to trunk

Group 0: <humanoid>(101) - Sub-goal: **Ensure** clear paths for mobile cars

Group 1: <mobile_car_2>(202) - Sub-goal: Transport the right wheel to franka assembly area

Group 2: <robot arm>(606) - Sub-goal: **Attach** the right wheel to the trunk

Figure 9: **Scene Dynamics Log**. It illustrates the formation of sub-groups and their corresponding sub-tasks.

Humanoid Prompt

Think of yourself as **humanoid (101)**, and your available skills are:

Available Skills for Humanoid (101):

- [walk] <humanoid> (101) move to selected area
- [carry] <humanoid> (101) carry <obstacles> (507)
- [wait] <humanoid> (101) wait

Your role: The humanoid robot has advanced bipedal locomotion and manipulation capabilities. You can walk around the assembly workspace, navigate between different areas, and handle obstacles. You specialize in mobility and can **carry obstacles (507)** to clear paths for other agents. You work collaboratively with wheeled robots and the franka arm to complete assembly tasks.

Note: To perform actions, you must **first [walk]** to the target area. For example, to carry obstacles, you must first walk to the area where the obstacles are located.

Note: Your primary function is obstacle management using the **[carry]** skill for <obstacles> (507). You work collaboratively with **wheeled robots (202, 203, 204)** for component transportation and with **franka (606)** for assembly operations.

Note: You can navigate complex terrain and access areas that wheeled robots might find difficult, making you valuable for clearing paths and removing obstructions.

Note: Use **[wait]** when you need to pause and coordinate with other agents or when your immediate assistance is not required.

Franka Prompt

Think of yourself as **franka (606)**, and your available skills are:

Available Skills for Franka (606):

- [check] <franka> (606) check <trunk> (303)
- [check] <franka> (606) check <left wheel> (405)
- [check] <franka> (606) check <right wheel> (406)
- [pick] <franka> (606) pick and place <left wheel> (405) on <trunk> (303)
- [pick] <franka> (606) pick and place <right wheel> (406) on <trunk> (303)
- [wait] <franka> (606) wait

Your role: You are a fixed manipulator robot arm specialized in precise assembly operations. You can check components and perform pick-and-place operations to assemble the final robot. You work with components that are brought to your workspace by **wheeled robots (202, 203, 204)**. Your primary task is to assemble wheels onto the trunk to complete the robot assembly.

AGV Prompt

Think of yourself as one of the **wheeled robots (202, 203, or 204)**, and your available skills are:

Available Skills for Wheeled Robots:

- [move] <wheeled robot1/2/3> (202/203/204) move to component location using RRT path
- [push] <wheeled robot1/2/3> (202/203/204) push selected component to franka area
- [wait] <wheeled robot1/2/3> (202/203/204) wait

Your role: You are a wheeled robot optimized for efficient component transportation around the assembly workspace. You can navigate to component locations using RRT path planning and push components like <trunk> (303), <left wheel> (405), and <right wheel> (406) to the franka assembly area. You work collaboratively with other wheeled robots, the humanoid, and the franka arm.

Note: Normally you must **first use [move]** skill to navigate to a component's location **before you can [push]** it. However, if your available actions only include [push] and [wait] (no [move] option), this means you have already completed the movement phase and are positioned at the component location. In this case, you can directly use the [push] skill to transport the component to the franka assembly area.

Note: You specialize in ground-level component transportation using RRT path planning. You cannot perform assembly operations - those require the franka (606) robot arm's assistance.

Note: When transporting components, coordinate with other wheeled robots to efficiently move all required components. **Prioritize moving the <trunk> (303) first**, then the wheels <left wheel> (405) and <right wheel> (406).

Note: Your **[push]** action transports components directly to the franka assembly area where the robot arm can access them for final assembly operations.

Note: Use **[wait]** when you need to pause for coordination with other agents or when your immediate assistance is not required.

Figure 10: **Single Agent Prompts.** It illustrates the single-agent prompt design for Franka, Humanoid, and AGV, specifying their roles and available skills.

Agent Grouping Vanilla Prompt

You are an intelligent task coordinator for multi-agent assembly systems. Your job is to analyze the current situation and develop a comprehensive strategy for grouping agents to accomplish the robot assembly task efficiently.

Given the following information:

These robot agents and their **capabilities and action spaces** are:

1. **humanoid (101)**: The humanoid can walk around the environment and carry obstacles to clear paths for other robots. The humanoid has advanced mobility and can navigate complex terrain. When other agents cannot reach component locations due to obstacles, humanoids can assist by clearing the path. The humanoid can carry obstacles but cannot perform component assembly directly.
NOTE: The humanoid should focus on path clearing and obstacle removal tasks.

2. **mobile_car (201/202/203)**: The mobile cars are wheeled robots optimized for component transportation. Each mobile car is assigned to a specific component:
- **mobile_car_1 (201)**: Responsible for left wheel transportation
- **mobile_car_2 (202)**: Responsible for right wheel transportation
- **mobile_car_3 (203)**: Responsible for trunk transportation
The mobile cars can navigate to component locations using RRT path planning and transport components to the franka assembly area.
Mobile cars cannot perform assembly operations directly but are essential for component logistics.

3. **franka robot arm (606)**: The franka is a fixed manipulator specialized for precise assembly operations. The franka can check components, pick and place wheels on the trunk to complete the final assembly. The franka can check trunk (303), left wheel (405), and right wheel (406). The franka performs the final assembly by placing wheels on the trunk. The franka cannot reach ground-level objects directly and relies on mobile cars to bring components to its workspace.

Available Agents: **#AGENTS_INFO#**

Task Goal: **#TASK_GOAL#**

Current Environment Observations: **#OBSERVATIONS#**

Previous Dialogue History: **#DIALOGUE_HISTORY#**

Pay attention to the dialogue history, because it records your conversations with different agents and the progress of the task execution. Depending on the content of the dialogue, you can avoid repeating actions that have been completed and make more informed grouping decisions based on what has already been accomplished.

Please develop a comprehensive strategy for agent grouping and task coordination. Consider the following aspects:

1. ****Situation Analysis****: Analyze the current environment state, what has been accomplished (based on dialogue history), and what still needs to be done.
2. ****Spatial Analysis****: Consider the physical locations of agents and components, area layouts, and any spatial constraints that might affect agent coordination.
3. ****Task Decomposition****: Break down the overall assembly task into logical sub-tasks that can be assigned to different groups.
4. ****Grouping Strategy****: Explain your reasoning for how to group agents, considering:
 - Which agents can work **simultaneously** without interfering with each other
 - Which agents need to **collaborate closely** and should be in the same group
 - Spatial separation to avoid conflicts
 - Efficiency considerations
 - Remember that agents going to interact with the same component and in the same area should be in the same group.
5. ****Sub-goal Assignment****: For each proposed group, describe:
 - What specific sub-goal they should accomplish
 - Why this sub-goal makes sense for this particular combination of agents
 - How this sub-goal contributes to the overall assembly task completion
6. ****Coordination Strategy****: Explain how the different groups will coordinate with each other, if needed, and the sequence or timing of their operations.
7. ****Risk Assessment****: Identify potential issues or conflicts that might arise and how your grouping strategy addresses them.

Please provide a detailed explanation on Grouping Strategy and Sub-goal Assignment part of your strategy while keeping the other parts brief. Focus on the reasoning and rationale rather than strict formatting. Be comprehensive in your analysis and make sure to address all aspects of effective multi-agent coordination.

Let's think step by step and develop the best possible strategy for this assembly situation.

Figure 11: **Vanilla Agent Grouping Prompt**. It illustrates the prompt to extract agent sub-groups and their corresponding sub-tasks with the general proposal planner.

Agent Grouping Prompt

You are a structured output formatter for multi-agent task coordination. Your job is to convert a detailed grouping strategy into a precise, parseable format.

Based on the comprehensive grouping strategy provided below, please extract and format the agent groups and their sub-goals according to the exact format requirements.

Grouping Strategy: #VANILLA_GROUPING_OUTPUT#

Available Agents: #AGENTS_INFO#

IMPORTANT FORMAT RULES:

1. Each group must be on a single line starting with "Group X:" where X is a number
2. List agents using EXACTLY this format: <agent_class_name>(id)
3. Separate multiple agents with commas
4. After listing agents, use " - Sub-goal: " followed by a clear, specific sub-goal
5. Sub-goals must be actionable and directly related to the task
6. Do not include any extra text, comments or explanations

Example correct format:

Group 1: <humanoid>(101), <mobile_car_1>(201) - Sub-goal: Transport trunk from area A to franka assembly area
Group 2: <franka>(606) - Sub-goal: Assemble left wheel onto trunk when components are ready

For agents that should not participate in this step:

Non-assigned Agent: <mobile_car_2>(202) - Reason: No tasks available at current location

INVALID FORMATS (DO NOT USE):

- Group 1 - The humanoid and mobile car will...
- First group: humanoid(101), mobile_car_1(201)
- Group 1: <humanoid>(101) will walk while <mobile_car_1>(201) moves
- **Group 1:** <humanoid>(101)
- Group 1: <humanoid>(101) - The agent should move to...

Remember:

- Use exact agent class names and IDs from the available agents list
- Keep groups spatially separated to avoid conflicts
- Only include agents that are actually needed for the current step
- Agents from all groups are expected to operate concurrently, so group them together if they need to act sequentially instead of act in parallel
- Provide clear, actionable sub-goals for each group
- Never assign the same agent to multiple groups
- Never include extra comments or explanations
- If previous plan failed, reflect and improve your plan by assigning sub-tasks to proper agent groups
- If part of the plan is already done, you don't need to assign that part of task again.
- Group agents together if they share similar subgoals (e.g. interacting with the same object)

Let's think step by step.

Figure 12: **Agent Grouping Prompt.** It illustrates the prompt to form agent sub-groups with a general proposal with the general proposal planner.