2026 IEEE International Conference on Robotics and Automation (ICRA 2026)
June 1-5, 2026. Vienna, Austria

Emergent Co-Adaptive Strategies in Heterogeneous Multi-Robot Systems via Meta-Learning

Haocheng Wang[1,2], Lin Wang[1], Tin Lun Lam[1,2], Jianwang Zhai[1], Dahai Lin[3], Hexiang Zheng[3], Xvchun He[1], Yuan Gao[1,2†]

***Abstract*— As teamed robots increasingly share public spaces with humans, the ability to co-adapt—to mutually adjust behavior in response to one another—becomes essential for safe, efficient, and socially acceptable operation. This paper introduces a socially co-adaptive framework for heterogeneous multi-robot systems (HMRS) that enables real-time adaptation to human behavior while preserving cooperative task execution. Our approach fuses large language models for natural language understanding with model-agnostic meta-learning to allow robots to rapidly generalize across diverse social contexts. We implement and validate the system using a real-world HMRS composed of robots with different roles—workers, a station, and a social robot—interacting with 44 human participants under induced behavioral states (relaxed vs. nervous). Results reveal significant behavioral adaptation: the system dynamically shifts between egoistic and altruistic strategies, improving crowd guidance success by 21%. It also reduces human cognitive load—specifically, physical demands by 39% and temporal demands by 39%—while increasing trust by 16% and perceived anthropomorphism by 21%. This work demonstrates the feasibility of human-robot co-adaptation at scale, laying the groundwork for socially intelligent robotic systems capable of thriving in complex, human-centered environments.**

I. INTRODUCTION

The next century could witness the emergence of a planetary-scale robot civilization—an autonomous, self-organizing stratum of machines that co-evolves with human society much as multicellular life once reconfigured the biosphere. In this vision, billions of heterogeneous robotic agents—terrestrial, aquatic, and aerial—form a parallel ecology whose collective metabolism is powered by renewable energy and whose evolutionary dynamics are driven by in-situ fabrication, modular reconfiguration, and algorithmic mutation. These agents would not merely serve humanity; they would participate as co-authors in the production of knowledge, infrastructure, and even culture, negotiating norms through distributed protocols and blockchain-verified social contracts.

Achieving such long-term visions requires progress on more immediate challenges. At present, heterogeneous multi-robot systems (HMRS) have already been increasingly deployed across a wide range of application domains, including search and rescue, environmental monitoring, intelligent manufacturing, and urban services, where the complementary strengths of diverse robots enhance task efficiency and system robustness [1]. Particularly in socially interactive environments involving human–robot co-existence, HMRS hold the potential to reduce human cognitive load, improve service experience, and enhance environmental adaptability [2].

Shenzhen Institute of Artificial Intelligence and Robotics for Society, Shenzhen, China[1], School of Science and Engineering, The Chinese University of Hong Kong, Shenzhen[2], Shenzhen Smart City Technology Development Group Co., Ltd.[3]. This work is supported by the National Key R&D Program of China under Grant No. 2024YFB4505500 & 2024YFB4505503.

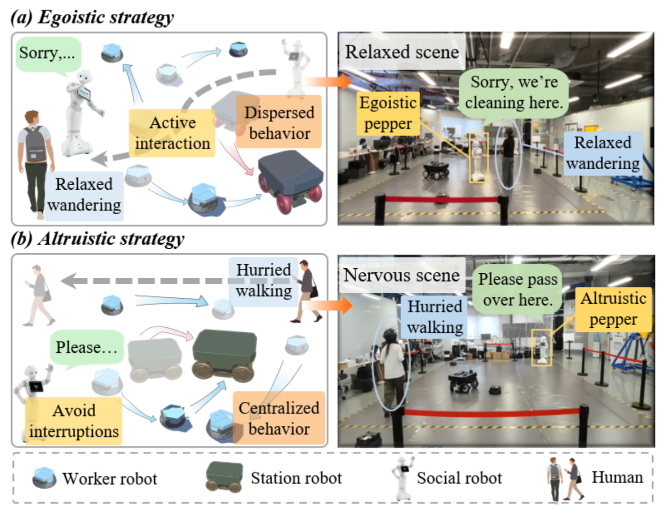

Fig. 1: Behavioral adaptations of the heterogeneous multi-robot system under the co-adaptation framework. Workers perform area coverage, stations provide energy replenishment, and pepper engages with humans as a social robot. (a) The system adapts to relaxed crowds with an egoistic strategy, where pepper interacts actively and other robots perform dispersed coverage tasks. (b) The system adapts to nervous crowds with an altruistic strategy, where pepper avoids interruptions and other robots adopt centralized behaviors. Right panels show corresponding real-world experimental scenes validating the framework.

However, existing HMRS often lack social adaptability, which hinders their ability to interact seamlessly with humans in dynamic and complex environments [3]. Such limitations can cause misalignment with human movement patterns, disruption of activity rhythms, and even a decline of user trust [4]. Prior studies further show that robotic systems without sufficient social adaptability struggle to integrate into human social contexts, thereby undermining their acceptability and long-term effectiveness [5].

To address this issue, we propose establishing a co-adaptation process between robotic systems and society. In general, co-adaptation refers to a process where agents mutually influence and adjust their behaviors in coordination to accommodate the changes induced by the adaptive actions of their counterparts [6]. In robotics, the concept is interpreted differently across subfields: within multi-robot systems, it emphasizes the ability of agents to dynamically coordinate and respond to one another, whereas in social robotics, it further entails perceiving and adapting to human behavior, as well as fostering mutual trust between humans and robots [7].

Building on the concept of co-adaptation, we aim to endow HMRS with the ability to dynamically align with human social behavior patterns while maintaining adaptability in multi-robot collaboration, as illustrated in Fig. 1. Although prior studies have explored co-adaptation in single-robot systems [8], human–robot teamwork [9], and social robot interactions [10], systematic investigations of co-adaptation in HMRS under social contexts remain scarce. This gap limits their deployment in public environments, mixed human–robot teams, and large-scale societal applications.

The goal of this study is to evaluate the adaptive performance of HMRS in social scenarios and assess their potential for seamless integration into human society. Key challenges include: (**1**) how to effectively simulate realistic social contexts in experiments; (**2**) how to characterize the behavioral changes of HMRS during interaction; and (**3**) how to capture both subjective and objective human responses through psychological and behavioral measurements.

To address these challenges, we developed an HMRS prototype with multi-robot collaboration and human–robot interaction capabilities, deployed in a controlled experimental setting. Through carefully designed tasks, the environment was configured as a socially rich interactive space. During the experiments, we simultaneously measured both the adaptive behaviors of the system and the psychological responses of participants, enabling an evaluation of the effectiveness and acceptability of HMRS adaptability at the social level.

The main contributions of this work are as follows:

- We formulate *social co-adaptation* as an online meta-learning objective for HMRS for the first time, embedding both crowd convergence theory and an egoism–altruism spectrum into the reward to enable few-shot policy switching while keeping HMRS' task performance.
- We introduce a language–behavior dual-channel adaptation architecture: an LLM continuously estimates the crowd's psychological urgency, while the meta-policy's egoism weight is adjusted in real time based on observed human trajectories. Together, they drive heterogeneous robots to reshape formations and utterances—without retraining. Real-world simulated airport-flow experiments with 44 participants show the robot raises crowd-guidance success by 21% within 3 minutes while cutting human physical & temporal load by 39%, yielding the first quantitative closed-loop where *both* humans and HMRS adapt to each other—prior work only achieved unilateral adaptation or avoidance.

II. Related Work

A. Heterogeneous Multi-Robot Systems

HMRS exploit the complementary capabilities of diverse platforms, offering advantages in efficiency and versatility [1]. Recent work has proposed task allocation, communication, and strategy optimization methods to address the coordination challenges arising from structural and functional heterogeneity, thereby enhancing stability and scalability [11], [12].

Despite this progress, several challenges remain. First, heterogeneity leads to discrepancies in sensing modalities and functional capabilities, leaving multi-source data fusion and consistency unresolved [13], [14]. Second, in dynamic environments, real-time perception and decision-making are still inadequate [15], [16]. Third, most studies emphasize technical feasibility but pay limited attention to human behavioral and social factors. Consequently, while HMRS have advanced technically, their adaptability and acceptability in social contexts remain underexplored. Our study addresses this gap by proposing and validating a socially co-adaptive framework to enhance both adaptability and social integration of HMRS in complex human–robot interaction settings.

B. Social Adaptability in Robotic Systems

Extensive work has examined social adaptability in single-robot tasks. In social robotics, adaptive adjustments are considered essential for human interaction, such as modifying dialogue strategies or behaviors based on user states [4], [17]. Another line of research focuses on enabling robots to navigate crowded environments while conforming to social norms [18], [19]. Other studies highlight challenges of deploying robots in public spaces, including security, privacy, and social acceptance [20], [21].

In multi-robot systems, researchers have begun incorporating human state awareness and behavioral preference modeling to enhance adaptability [22], [23]. However, most

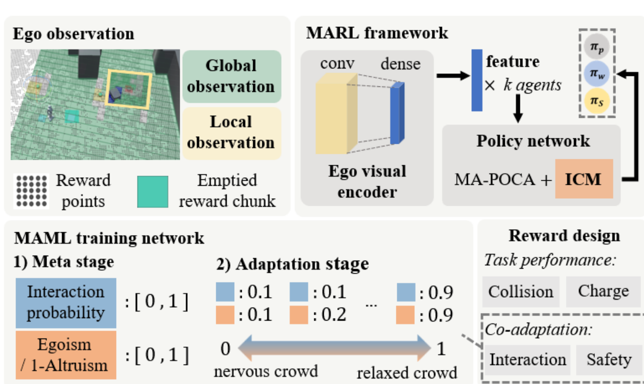

Fig. 2: Policy structure of the proposed framework. It consists of three core components: (i) ego observation, which unifies perception across robot types; (ii) the MARL framework, which coordinates heterogeneous agents under a shared reward; and (iii) the MAML training network, which enables rapid policy adaptation to new social conditions.

efforts focus on homogeneous systems, with little systematic analysis of HMRS in social contexts. Moreover, current studies are largely limited to theoretical models and simulations [24], with few validated in real human–robot interactions. In contrast, our work develops an experimental platform that operates in real social scenarios, enabling empirical evaluation of HMRS social adaptability and providing practice-oriented insights for future research.

III. METHOD

A. Heterogeneous Multi-Robot System

Building on the worker-station coverage framework in [25], our HMRS is designed for tasks that combine area coverage, energy supply and human interaction. Workers collaborate to achieve efficient coverage of the target area, while stations provide real-time energy replenishment to sustain worker operation. To extend this setting toward socially adaptive contexts, we further introduce pepper as a social robot, responsible for engaging with humans and facilitating co-adaptation.

1) Agent Composition: We model five categories of agents, denoted as $\{W, S, P, A, D\}$, representing workers, stations, pepper robots, walkers, and dynamic obstacles. Walkers correspond to human participants in the environment, while dynamic obstacles denote non-human disturbances such as mobile carts. Agents are grouped into robots $R \equiv \{W, S, P\}$ and external disturbances $E \equiv \{A, D\}$. For instance,

$$W \equiv \{W^1, W^2, \ldots, W^{nw}\},$$

where n_W is the number of worker robots.

2) Team Reward Design: To align heterogeneous agents within group $\mathcal{R}$ toward cooperative behavior, we design a unified reward function that balances task efficiency, safety, and social interaction, defined as

$$r = \sum_{w \in W} (r_c^w + r_e^w) + \sum_{s \in S} r_m^s + \sum_{p \in P} r_i^p + r_t + r_{\text{collision}} \quad (1)$$

where r_c^w and r_e^w denote coverage and energy rewards for workers, r_m^s penalizes unnecessary station movement, and r_i^p captures pepper's interaction reward with walkers. The temporal reward r_t provides a constant bonus before task completion to encourage efficiency, while $r_{\text{collision}}$ penalizes collisions among all agents to promote safe operation. This shared formulation simplifies reward design across agent types and provides the foundation for the policy structure described in Section III-B.

B. Policy structure

The overall policy integrates three components—ego observation, the MARL framework, and the MAML training network—as illustrated in Fig. 2.

1) Ego observation: To unify heterogeneous perception, we adopt a dual-grid sensing model that provides consistent observations across robot types. Each robot constructs (i) a high-resolution local grid for navigation and obstacle avoidance, and (ii) a low-resolution global grid for human motion prediction. Further details are available in [26].

2) MARL framework: Policies are optimized under the shared reward function defined in Section III-A.2. Our system builds on the cooperative MA-POCA baseline [27], enhanced by three modules: (i) scaled dot-product attention for efficient observation aggregation in large-scale environments [28], (ii) an Intrinsic Curiosity Module (ICM) [29] to provide exploration incentives under sparse rewards, and (iii) curriculum learning [30] to stabilize training. As illustrated in Fig. 3, curriculum learning progresses through three stages: Stage I focuses on basic behavior learning, Stage II emphasizes collaboration, and Stage III integrates human factors for co-adaptive interaction.

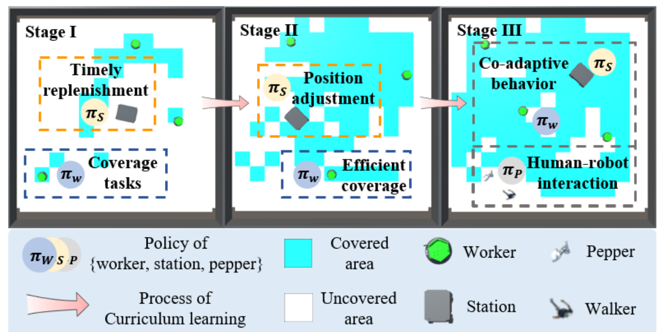

Fig. 3: The curriculum is divided into three stages. In Stage I, the worker and station robots learn basic behaviors: workers identify uncovered areas and complete coverage tasks, while stations follow workers and provide timely replenishment. Stage II focuses on collaboration: workers autonomously divide the workspace, while stations adjust their positions for non-intrusive support. In Stage III, pepper and walker agents are introduced, enabling human–robot interaction and training of co-adaptive capabilities.

3) MAML training network: To further enhance adaptability, we adopt MAML [31], which meta-trains across auxiliary environments with varied human densities and behavioral states. The meta policy learns an initialization that, with a few gradient updates, adapts rapidly to new crowd conditions. Combined with curriculum learning and attention, this structure enables fast policy and critic adjustment under diverse environmental distributions.

C. Human-robot Co-adaptation

Our HMRS targets social environments composed of human groups, where group-level states influence human–robot interactions. We focus on two dimensions that capture such effects: *collective behavior characteristics* and the *egoism–altruism* spectrum, which are integrated into the meta-learning framework.

1) Collective Behavior Characteristics: Convergence theory suggests that crowds form through shared goals or interests, such as airport passengers rushing to catch a flight. Urgency thus affects their likelihood of engaging with robots, as validated in prior studies [32]. In our experiments, we approximate this engagement with a Gaussian model

$$p_{\text{interaction}} \sim \mathcal{N}(\mu_p, \sigma_p^2) \tag{2}$$

where μ_p denotes the average crowd interaction probability across crowds, which reflects how likely a crowd is to engage with the social robot, and σ_p^2 captures its variance. In practice μ_p is estimated from observed crowd signals (*e.g.*, proportion of crowd-initiated communications) and is used to condition the reward structure that guides policy adaptation.

2) Egoism and Altruism: We characterize social tendencies of the HMRS with two dimensions: egoism $\mathcal{E}$ and altruism $\mathcal{A}$. Egoistic systems intervene actively to minimize external disturbance, while altruistic systems reduce their own influence to maintain social tolerance. Formally, these attributes parameterize the policy model $B(\theta; \mathcal{E}, \mathcal{A})$, where θ represents the general policy parameters. The model output is a behavior policy (*i.e.*, action distribution) that trades off task objectives (*e.g.*, coverage, energy) with social objectives induced by $\mathcal{E}$ and $\mathcal{A}$.

3) Meta Learning Environment Setup: We define $\mathcal{H}$ as a high-dimensional task space, where each $h \in \mathcal{H}$ represents a human–robot interaction scenario. Using MAML, we optimize θ such that it can rapidly adapt to scenario-specific loss functions $L_{\mathcal{T}}$ drawn from a distribution $\mathcal{T} \sim \mathcal{H}$, formulated as

$$\min_{\theta} \mathbb{E}_{\mathcal{T} \sim \mathcal{H}} \left[L_{\mathcal{T}} \left(U(\theta, \mathcal{T}) \right) \right] \tag{3}$$

where $U(\theta, \mathcal{T})$ represents the policy parameters updated through task $\mathcal{T}$. Auxiliary tasks vary in human density and behavioral urgency, allowing the meta-policy to generalize across diverse social contexts.

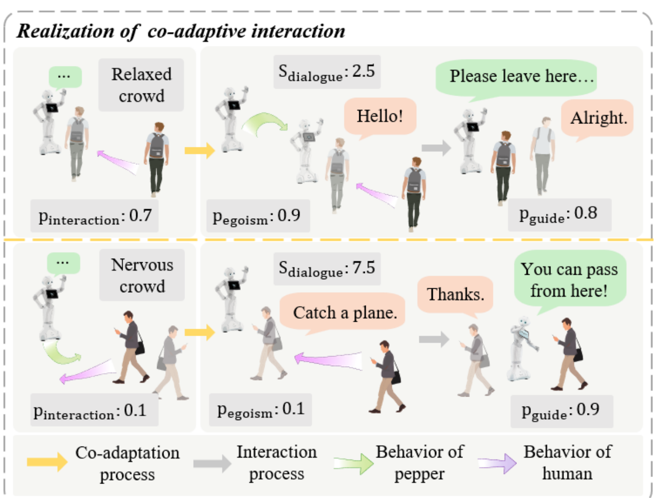

Fig. 4: Illustration of co-adaptive interaction. Four key indicators are visualized: (i) $P_{\text{interaction}}$ denotes the proportion of crowd-initiated interactions with pepper; (ii) P_{ego} represents the proportion of pepper's proactive interception attempts; (iii) P_{guide} indicates the success rate of pepper guiding the crowd to pass through; and (iv) S_{dialogue} reflects large language models' evaluation of human utterances. These measures jointly characterize the system's co-adaptive interaction dynamics.

To operationalize co-adaptation, we shape episode rewards using both group-level interaction $p_{\text{interaction}}$ and egoism level p_{ego}. Concretely,

$$r_{\text{interaction}} = \begin{cases} C \cdot p_{\text{ego}}, & \text{if } p < p_{\text{interaction}} \\ -C \cdot (1 - p_{\text{ego}}), & \text{otherwise} \end{cases} \tag{4}$$

where p is the observed engagement in the episode (*e.g.*, fraction of crowd responses), $p_{\text{ego}} \in [0, 1]$ parameterizes egoistic tendency, and $C > 0$ scales the social reward. Intuitively, when observed engagement p is below the group expectation $p_{\text{interaction}}$, we reward egoistic actions (*i.e.*, higher p_{ego}) to solicit interaction; otherwise, we penalize intrusive egoistic behaviors to encourage altruistic behavior. Embedding this reward into the meta-training tasks biases the learned initialization θ toward fast co-adaptive behavior across varying group states.

D. Realization of Co-adaptive Interaction

To implement the co-adaptive strategies, we simulate an airport environment where workers clean designated areas, a station provides energy replenishment, and a pepper robot guides and interacts with passengers. This setup enables the application of the co-adaptation framework in a controlled setting. Specifically, it enables the theoretical constructs of group-level behavioral traits to be expressed through measurable interaction patterns of the robots and humans, which are visualized in Fig. 4.

1) Definition of the Interaction Process: To capture co-adaptation in practice, we focus on two probabilistic parameters: $p_{\text{interaction}}$ and p_{ego}. $p_{\text{interaction}}$ represents the expected probability of crowd engagement with pepper, estimated from the ratio of crowd-initiated communications to total passages. A lower observed engagement relative to this expectation indicates weak responsiveness, prompting higher p_{ego} (*i.e.*, the ratio of pepper's proactive interaction attempts). Conversely, higher observed engagement allows the system to reduce interventions. Together, these parameters quantify behavioral tendencies and directly connect to the reward shaping described in Section III-C.3.

2) Behavioral State Induction: Based on the interaction process, nervous crowds are expected to exhibit lower $p_{\text{interaction}}$ and p_{ego}, while relaxed crowds show higher values. As reproducing real urgency is infeasible in the lab, so we induced participant behavioral states through imagined scenarios. Following a role-playing paradigm [33], participants envisioned either rushing to catch a flight (*i.e.*, nervous) or walking leisurely (*i.e.*, relaxed), with emotional immersion reinforced by images of airport terminals [34]. Such paradigms are widely adopted to induce behavioral states without requiring real travel.

3) Adaptive Dialogue: Beyond behavioral adaptation, pepper adjusted dialogue strategies to participants' psychological state. Pepper was equipped with LLMs for real-time interaction, as recent work shows they outperform traditional methods in emotion recognition and can directly assign affective scores to dialogue [35], [36]. Building on this, we

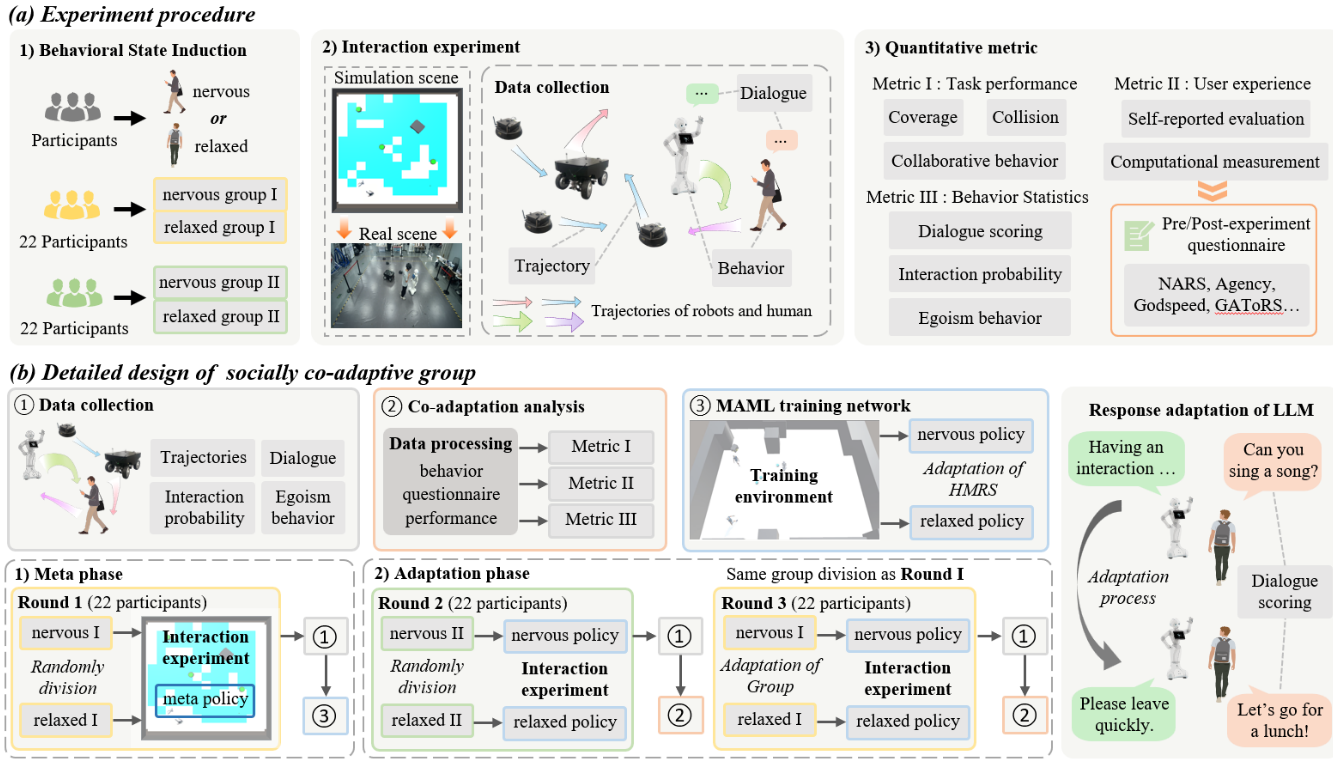

Fig. 5: Experimental design of the socially co-adaptive study. (a) Overall experimental procedure, including behavioral state induction, interaction with the HMRS, data collection, and multi-dimensional evaluation through three categories of metrics: task performance, behavioral statistics, and user experience. (b) Detailed design of the co-adaptive experiment: Round 1 establishes baseline data with a unified meta-policy; Round 2 evaluates state-specific adapted policies on new participants; Round 3 re-tests the initial participants to validate bidirectional human–robot co-adaptation.

designed a prompt-based scoring mechanism that conditions the LLMs on the role of a social robot. After each round of interaction, pepper estimated participants' urgency on a 1–10 scale (*i.e.*, 1 = completely relaxed, 10 = highly nervous). High urgency scores triggered more polite and concise directive utterances, whereas low scores elicited more direct and socially persuasive responses. The aggregated scores then guided pepper's dialogue strategy, enabling adaptive verbal interaction aligned with crowd states and further reinforcing bidirectional co-adaptation.

IV. Experiments & Results

A. Experimental Design

1) Experimental content: This experiment evaluated both task performance and human-robot co-adaptation in a simulated airport environment—a representative high-density social setting where urgency strongly shapes interaction. As illustrated in Fig. 5(a), participants were briefed on the scenario and assigned to either a nervous (*i.e.*, simulating hurried travelers) or relaxed (*i.e.*, simulating passengers with ample time) groups based on behavioral state induction. During the experiment, we collected trajectories, actions, and dialogues, followed by validated questionnaires assessing participants' psychological states and attitudes.

2) Evaluation metrics: We defined three categories of evaluation metrics. First, task performance metrics included coverage rate, collision incidents, and cooperative behaviors, ensuring that social adaptation did not compromise baseline functionality. Second, behavioral statistics metrics captured the system's co-adaptive characteristics, such as interaction probability, the proportion of egoistic actions, the success rate of guidance, and dialogue scores. Finally, user experience metrics measured participants' adaptation and acceptance through validated scales, covering workload (*e.g.*, NASA-TLX [37]), trust and attitudes (*e.g.*, Trust Scale [38], NARS [39], GAToRS [40], Agency Questionnaire [41], and perception of robots (*e.g.*, Godspeed [42]). Together, these metrics provided a comprehensive evaluation across technical, behavioral, and psychological dimensions.

3) Experimental procedure: As shown in Fig. 5(b), 44 volunteers were recruited and randomly assigned into nervous or relaxed groups (*i.e.*, 22 each) via behavioral state induction. In Round 1, 11 participants per group interacted with a unified meta-policy system, and their behavioral data were used to generate state-specific adaptive strategies. In Round 2, the remaining 11 participants per group were tested with the corresponding "nervous-adapted" or "relaxed-adapted" system to evaluate adaptation effects on new participants. In Round 3, the initial 22 participants returned, were re-induced into their original states, and interacted with the updated system. This design allowed us to examine both system adaptation to humans and human adaptation to the system. Behavioral and questionnaire data were analyzed using statistical tests (*i.e.*, t-tests, significance set at $p < 0.05$).

B. Experimental Setup

1) Environment: Physical experiments were conducted in a laboratory testbed that emulates an airport cleaning zone. A $6 \times 6\,\text{m}$ area was delineated using floor markings to indicate restricted regions, and physical safety barriers prevented accidental entry by non-participants. The layout mirrored the simulated environment to ensure geometry and task constraints were comparable.

2) Robot Team Setup: The heterogeneous team comprised three SPARK worker robots[1], one custom four-wheel-drive robot acting as the station, and one pepper robot as the social robot[2]. This composition matches the simulation setup to maintain consistency across domains.

3) Experimental Equipment Setup: A high-performance server (*i.e.*, Intel Core i9-10900K, 10 cores/20 threads, 32 GB RAM, Ubuntu 20.04 LTS) orchestrated the system. A motion-capture system tracked robot and human trajectories; a calibrated mapping aligned simulation and physical coordinate frames. Communication used a low-latency Wi-Fi router within a closed local network. All trained models and runtime modules were deployed as ROS nodes (ROS Noetic), handling inference, data logging, and inter-robot communication.

C. Result analysis

This study conducted 66 experiments. The three experimental rounds are denoted by N_1, N_2, N_3 for the nervous group and R_1, R_2, R_3 for the relaxed group. Each interaction session lasted three minutes. The recorded metrics were then subjected to statistical analysis.

1) Task performance: The real-world system was directly mapped from the simulation environment, where communication latency and hardware constraints could potentially introduce deviations. Statistical analysis showed that these effects were negligible.

Across 66 experiments, the system experienced 12 inter-robot collisions, none of which disrupted task execution. There were no instances of robot–human collisions, demonstrating that the system maintained adequate safety. In cooperative tasks, 84 energy-supply interactions between the station S and workers W were recorded, with only 4 failures. These failures were primarily due to network latency and did not affect overall task completion. For coverage, as shown in Fig. 6(c), worker robots consistently achieved full coverage in simulation, and nearly identical results ($M = .994, SD < .001$) in real-world deployment. These findings confirm that the HMRS maintained stable task performance in physical trials, with no significant degradation relative to simulation baselines.

[1] https://github.com/NXROBO/spark

[2] https://www.softbankrobotics.com/jp/product/pepper/

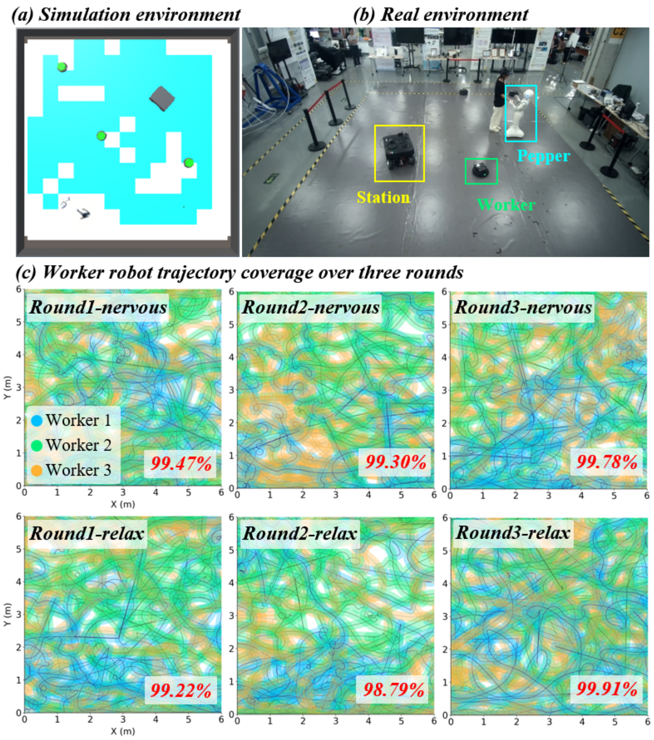

Fig. 6: (a) Simulated environment; (b) corresponding real-world experimental setup; (c) coverage rates of worker robots for both nervous and relaxed groups in each round. Results show that all trials achieved near-complete coverage ($M = .994, SD < .001$), closely matching the 100% coverage observed in simulation, indicating that the HMRS maintained task performance.

2) Behavioral Analysis: Behavioral state induction proved effective: the relaxed group showed a significantly higher interaction probability than the nervous group, $t(42) = 16.291, p < .001$, confirming that urgency was successfully manipulated. Before adaptation, egoistic tendencies were comparable across groups. After adaptation, however, distinct patterns emerged. Egoism increased markedly in the relaxed group R_2 (*i.e.*, $t(20) = -9.938, p < .001$) and R_3 (*i.e.*, $t(10) = -10.908, p < .001$) , while it decreased in the nervous group N_2 (*i.e.*, $t(20) = 15.871, p < .001$) and N_3 (*i.e.*, $t(10) = 18.338, p < .001$). This divergence indicates that the system learned to proactively engage relaxed participants while adopting more restrained, altruistic strategies toward nervous participants.

Guidance success also improved across rounds, reaching its highest level in R_3 (*i.e.*, $p_{\text{guide}} = .903$). This suggests that repeated exposure fostered bidirectional co-adaptation, with the system refining its guidance strategy while human participants increasingly aligned their behaviors, leading to higher success rates in Round 3.

In terms of linguistic interaction, dialogue scores assigned by the LLMs further distinguished the two groups: the nervous group received significantly lower scores than the relaxed group, $t(42) = -34.947, p < .001$, reflecting shorter and less socially engaging utterances under high urgency.

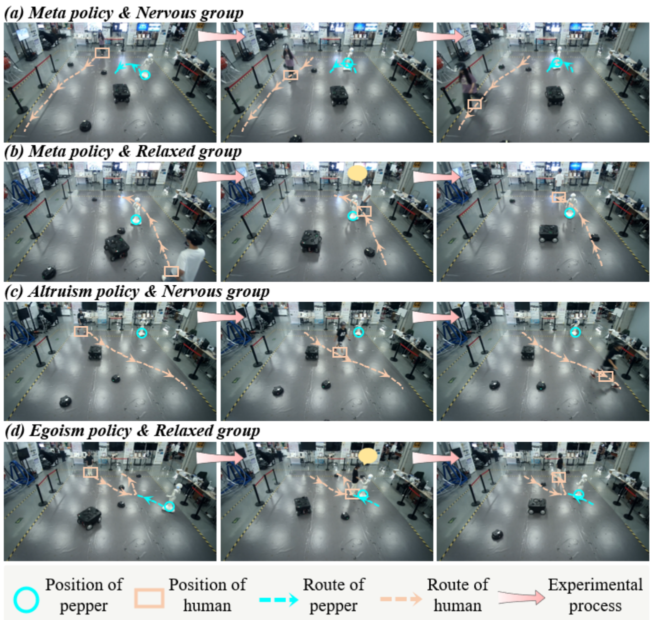

Fig. 7: Adaptive changes of the system across different stages. (a)-(b) Meta-policy learning of behavioral states for nervous and relaxed crowds, respectively. (c) After adaptation to nervous crowds, pepper adopts altruistic strategies (actively avoiding pedestrians).(d) After adaptation to relaxed crowds, pepper adopts egoistic strategies (actively intercepting and guiding).

The specific values of all behavioral indicators are listed in Table I.

As illustrated in Fig. 7, these results collectively demonstrate co-adaptation at the behavioral level: the HMRS dynamically shifted its strategy according to crowd state, while human participants reciprocally adapted to the system's interventions.

TABLE I: Mean values of behavioral indicators—interaction probability, egoistic tendency, guidance success, and dialogue scores—across different rounds(*i.e.*, N_1–N_3, R_1–R_3) in real-world experiments.

Experiment	Group	$P_{\text{interaction}}$	P_{ego}	P_{guide}	S_{dialogue}
Round 1	N_1	0.133	0.500	0.975	7.755
	R_1	0.491	0.506	0.653	2.431
Round 2	N_2	0.080	0.121	0.971	8.130
	R_2	0.548	0.845	0.745	2.664
Round 3	N_3	0.109	0.120	0.995	7.942
	R_3	0.488	0.863	0.903	2.796

3) User Experience: Questionnaire results showed that system adaptation significantly influenced participants' subjective experience. NASA-TLX scores revealed reduced workload in both groups: physical demands decreased (*i.e.*, N_2, $t(20) = 3.28, p < .05$; N_3, $t(10) = 4.26, p < .001$; R_2, $t(20) = 2.35, p < .05$; R_3, $t(10) = 1.99, p = 0.071$) and temporal demands also dropped (*i.e.*, N_2, $t(20) = 5.25, p < .001$; N_3, $t(10) = 5.37, p < .001$; R_2, $t(20) = 2.58, p < .05$; R_3, $t(10) = 2.57, p < .05$). These decreases indicate that dynamic adaptation reduced subjective burden. Godspeed results further showed higher ratings of anthropomorphism and animacy in R_2 (*i.e.*, anthropomorphism, $t(20) = -2.11, p < .05$; animacy, $t(20) = -4.24, p < .001$) and R_3 (*i.e.*, anthropomorphism, $t(10) = -2.55, p < .05$; animacy, $t(10) = -2.91, p < .05$) than in R_1, suggesting improved impressions of the robots.

Several other scales also reflected adaptive changes on the human side. In the third round, the 40-item trust scale indicated increased trust compared to earlier rounds (*i.e.*, N_1, $t(10) = 2.69, p < .05$; N_2, $t(20) = 3.31, p < 0.05$; R_1, $t(10) = 2.24, p < .05$; R_2, $t(20) = 2.48, p < .05$). Consistent with this, Agency and GAToRS scores revealed that nervous participants in Round 3 perceived stronger "thoughts" (*i.e.*, N_1, $t(10) = 2.71, p < .05$; N_2, $t(20) = 2.39, p < .05$) and social expectations (*i.e.*, N_1, $t(10) = 3.91, p < .001$; N_2, $t(20) = 2.85, p < .05$) from the system than former rounds, reflecting a shift in how humans regarded the robots as intentional social actors. Interestingly, NARS showed that the nervous group reported lower negative attitudes than the relaxed group in the interaction dimension (*i.e.*, $t(42) = -5.68, p < .05$) and social-influence dimension (*i.e.*, $t(42) = -5.72, p < .05$), suggesting that crowd state also shaped psychological judgments. Taken together, these psychological measures align with the behavioral findings: co-adaptation extended beyond observable behaviors to participants' workload, trust, and social attitudes, providing convergent evidence of bidirectional adaptation between humans and the HMRS. Detailed scale results are listed in Table II.

TABLE II: Mean values of participants' psychological measures on dimensions showing significant differences across validated scales.

Scale	Dimension	N_1	R_1	N_2	R_2	N_3	R_3
NASA-TLX	Physical demands	4.909	3.273	3.091	1.909	2.818	2.197
	Temporal demands	5.182	3.909	3.318	2.364	3.182	2.273
Godspeed	Anthropomorphism	3.428	3.273	3.455	3.873	3.691	4.018
	Animacy	3.572	2.848	3.424	3.673	3.618	3.782
40-item	Trust	6.018	6.073	5.582	5.836	7.495	7.545
Agency	Thoughts	3.182	3.636	3.400	3.758	4.055	3.430
GAToRS	Hopes at societal level	4.882	5.073	5.255	5.818	6.073	4.964
NARS	Interaction	2.455	3.127	2.652	3.578	2.564	3.455
	Social	2.766	3.424	2.691	3.576	2.982	3.788

V. Conclusion

This paper introduced a socially co-adaptive framework for HMRS that combines LLM-based intent understanding with meta-learning for rapid adaptation. We developed a prototype and validated it through real-world human–robot interaction experiments. Across task, behavioral, and psychological dimensions, the system showed stable deployment, bidirectional adaptation with humans, and improved human trust and social acceptance. These findings demonstrate the feasibility of socially co-adaptive HMRS and provide a foundation for their deployment in large-scale social environments.

REFERENCES

[1] Y. Rizk, M. Awad, and E. W. Tunstel, "Cooperative heterogeneous multi-robot systems: A survey," *ACM Computing Surveys (CSUR)*, vol. 52, no. 2, pp. 1–31, 2019.

[2] A. Dahiya, A. M. Aroyo, K. Dautenhahn, and S. L. Smith, "A survey of multi-agent human–robot interaction systems," *Robotics and Autonomous Systems*, vol. 161, p. 104335, 2023.

[3] L. Dong, Z. He, C. Song, X. Yuan, and H. Zhang, "Multi-robot social-aware cooperative planning in pedestrian environments using attention-based actor-critic," *Artificial Intelligence Review*, vol. 57, no. 4, p. 108, 2024.

[4] R. Jahanmahin, S. Masoud, J. Rickli, and A. Djuric, "Human-robot interactions in manufacturing: A survey of human behavior modeling," *Robotics and Computer-Integrated Manufacturing*, vol. 78, p. 102404, 2022.

[5] A. Noormohammadi-Asl, S. L. Smith, and K. Dautenhahn, "To lead or to follow? adaptive robot task planning in human-robot collaboration," *IEEE Transactions on Robotics*, 2025.

[6] S. Ahlberg, A. Axelsson, P. Yu, W. S. Cortez, Y. Gao, A. Ghadirzadeh, G. Castellano, D. Kragic, G. Skantze, and D. V. Dimarogonas, "Co-adaptive human–robot cooperation: summary and challenges," *Unmanned Systems*, vol. 10, no. 02, pp. 187–203, 2022.

[7] B. Gebru, L. Zeleke, D. Blankson, M. Nabil, S. Nateghi, A. Homaifar, and E. Tunstel, "A review on human–machine trust evaluation: Human-centric and machine-centric perspectives," *IEEE Transactions on Human-Machine Systems*, vol. 52, no. 5, pp. 952–962, 2022.

[8] A. Nanavati, V. Ranganeni, and M. Cakmak, "Physically assistive robots: A systematic review of mobile and manipulator robots that physically assist people with disabilities," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 7, 2023.

[9] H. ElMaraghy and W. ElMaraghy, "Adaptive cognitive manufacturing system (acms)–a new paradigm," *International Journal of Production Research*, vol. 60, no. 24, pp. 7436–7449, 2022.

[10] A. Kubota and L. D. Riek, "Methods for robot behavior adaptation for cognitive neurorehabilitation," *Annual review of control, robotics, and autonomous systems*, vol. 5, no. 1, pp. 109–135, 2022.

[11] L. Antonyshyn, J. Silveira, S. Givigi, and J. Marshall, "Multiple mobile robot task and motion planning: A survey," *ACM Computing Surveys*, vol. 55, no. 10, pp. 1–35, 2023.

[12] Z. Mandi, S. Jain, and S. Song, "Roco: Dialectic multi-robot collaboration with large language models," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 286–299, IEEE, 2024.

[13] X. Liu, S. Wen, J. Zhao, T. Z. Qiu, and H. Zhang, "Edge-assisted multi-robot visual-inertial slam with efficient communication," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 2186–2198, 2024.

[14] G. Wang, C. Zhang, S. Liu, Y. Zhao, Y. Zhang, and L. Wang, "Multi-robot collaborative manufacturing driven by digital twins: advancements, challenges, and future directions," *Journal of Manufacturing Systems*, vol. 82, pp. 333–361, 2025.

[15] C. Li, R. Zhang, J. Wong, C. Gokmen, S. Srivastava, R. Martín-Martín, C. Wang, G. Levine, M. Lingelbach, J. Sun, *et al.*, "Behavior-1k: A benchmark for embodied ai with 1,000 everyday activities and realistic simulation," in *Conference on Robot Learning*, pp. 80–93, PMLR, 2023.

[16] Y. Liu, W. Chen, Y. Bai, X. Liang, G. Li, W. Gao, and L. Lin, "Aligning cyber space with physical world: A comprehensive survey on embodied ai," *IEEE/ASME Transactions on Mechatronics*, 2025.

[17] K. Dautenhahn, "Socially intelligent robots: dimensions of human–robot interaction," *Philosophical transactions of the royal society B: Biological sciences*, vol. 362, no. 1480, pp. 679–704, 2007.

[18] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core challenges of social robot navigation: A survey," *ACM Transactions on Human-Robot Interaction*, vol. 12, no. 3, pp. 1–39, 2023.

[19] C. Chen, Y. Liu, S. Kreiss, and A. Alahi, "Crowd-robot interaction: Crowd-aware robot navigation with attention-based deep reinforcement learning," in *2019 international conference on robotics and automation (ICRA)*, pp. 6015–6022, IEEE, 2019.

[20] S. O. Oruma, M. Sánchez-Gordón, R. Colomo-Palacios, V. Gkioulos, and J. K. Hansen, "A systematic review on social robots in public spaces: threat landscape and attack surface," *Computers*, vol. 11, no. 12, p. 181, 2022.

[21] S. O. Oruma and S. Petrovic, "Security threats to 5g networks for social robots in public spaces: A survey," *IEEE Access*, vol. 11, pp. 63205–63237, 2023.

[22] C. Huang, W. Luo, and R. Liu, "Meta preference learning for fast user adaptation in human-supervisory multi-robot deployments," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5851–5856, IEEE, 2021.

[23] W. Wang, A. Bera, and B.-C. Min, "Hypergraph-based coordinated task allocation and socially-aware navigation for multi-robot systems," *arXiv preprint arXiv:2409.11561*, 2024.

[24] M. Viana, P. Alencar, E. Guimarães, E. Cirilo, and C. Lucena, "Creating a modeling language based on a new metamodel for adaptive normative software agents," *IEEE Access*, vol. 10, pp. 13974–13996, 2022.

[25] J. Tang, Y. Gao, and T. L. Lam, "Learning to coordinate for a worker-station multi-robot system in planar coverage tasks," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 12315–12322, 2022.

[26] Y. Gao, J. Chen, X. Chen, C. Wang, J. Hu, F. Deng, and T. L. Lam, "Asymmetric self-play-enabled intelligent heterogeneous multirobot catching system using deep multiagent reinforcement learning," *IEEE Transactions on Robotics*, vol. 39, no. 4, pp. 2603–2622, 2023.

[27] A. Cohen, E. Teng, V.-P. Berges, R.-P. Dong, H. Henry, M. Mattar, A. Zook, and S. Ganguly, "On the use and misuse of absorbing states in multi-agent reinforcement learning," *arXiv preprint arXiv:2111.05992*, 2021.

[28] Z. Shen, M. Zhang, H. Zhao, S. Yi, and H. Li, "Efficient attention: Attention with linear complexities," in *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 3531–3539, 2021.

[29] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, "Curiosity-driven exploration by self-supervised prediction," in *International conference on machine learning*, pp. 2778–2787, PMLR, 2017.

[30] Y. Bengio, J. Louradour, R. Collobert, and J. Weston, "Curriculum learning," in *Proceedings of the 26th annual international conference on machine learning*, pp. 41–48, 2009.

[31] E. Mitchell, R. Rafailov, X. B. Peng, S. Levine, and C. Finn, "Offline meta-reinforcement learning with advantage weighting," in *International Conference on Machine Learning*, pp. 7780–7791, PMLR, 2021.

[32] H. Y. Im, D. N. Albohn, T. G. Steiner, C. A. Cushing, R. B. Adams Jr, and K. Kveraga, "Differential hemispheric and visual stream contributions to ensemble coding of crowd emotion," *Nature human behaviour*, vol. 1, no. 11, pp. 828–842, 2017.

[33] L. Benson III and L. R. Beach, "The effects of time constraints on the prechoice screening of decision options," *Organizational Behavior and Human Decision Processes*, vol. 67, no. 2, pp. 222–228, 1996.

[34] J. J. Gross, "Emotion regulation: Current status and future prospects," *Psychological inquiry*, vol. 26, no. 1, pp. 1–26, 2015.

[35] H. Touvron, L. Martin, K. Stone, P. Albert, A. Almahairi, Y. Babaei, N. Bashlykov, S. Batra, P. Bhargava, S. Bhosale, *et al.*, "Llama 2: Open foundation and fine-tuned chat models," *arXiv preprint arXiv:2307.09288*, 2023.

[36] Y. Zhang, M. Wang, Y. Wu, P. Tiwari, Q. Li, B. Wang, and J. Qin, "Dialoguellm: Context and emotion knowledge-tuned large language models for emotion recognition in conversations," *Neural Networks*, p. 107901, 2025.

[37] S. G. Hart and L. E. Staveland, "Development of nasa-tlx (task load index): Results of empirical and theoretical research," in *Advances in psychology*, vol. 52, pp. 139–183, Elsevier, 1988.

[38] Y. Rong, T. Leemann, T.-T. Nguyen, L. Fiedler, P. Qian, V. Unhelkar, T. Seidel, G. Kasneci, and E. Kasneci, "Towards human-centered explainable ai: A survey of user studies for model explanations," *IEEE transactions on pattern analysis and machine intelligence*, vol. 46, no. 4, pp. 2104–2122, 2023.

[39] T. Nomura, T. Kanda, and T. Suzuki, "Experimental investigation into influence of negative attitudes toward robots on human–robot interaction," *Ai & Society*, vol. 20, no. 2, pp. 138–150, 2006.

[40] M. Koverola, A. Kunnari, J. Sundvall, and M. Laakasuo, "General attitudes towards robots scale (gators),"

[41] W. Wen and H. Imamizu, "The sense of agency in perception, behaviour and human–machine interactions," *Nature Reviews Psychology*, vol. 1, no. 4, pp. 211–222, 2022.

[42] C. Bartneck, D. Kulić, E. Croft, and S. Zoghbi, "Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots," *International journal of social robotics*, vol. 1, no. 1, pp. 71–81, 2009.