2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

Towards Human-Centered Construction Robotics: A Reinforcement Learning-Driven Companion Robot for Contextually Assisting Carpentry Workers

Yuning Wu[1], Jiaying Wei[1], Jean Oh[1], and Daniel Cardoso Llach[1]

Abstract— **In the dynamic construction industry, traditional robotic integration has primarily focused on automating specific tasks, often overlooking the complexity and variability of human aspects in construction workflows. This paper introduces a human-centered approach with a "work companion rover" designed to assist construction workers within their existing practices, aiming to enhance safety and workflow fluency while respecting construction labor's skilled nature. We conduct an in-depth study on deploying a robotic system in carpentry formwork, showcasing a prototype that emphasizes mobility, safety, and comfortable worker-robot collaboration in dynamic environments through a contextual Reinforcement Learning (RL)-driven modular framework. Our research advances robotic applications in construction, advocating for collaborative models where adaptive robots support rather than replace humans and underscores the potential for an interactive and collaborative human-robot workforce.**

I. INTRODUCTION

Construction remains the world's most labor-intensive industry [1], characterized by its manual nature and bespoke building requirements. Despite advancements like prefabrication [2] and on-site 3D printing [3], most construction tasks still depend heavily on skilled labor, operating in complex, dynamic, and cluttered environments. Workers not only endure significant physical strain but also perform frequent on-the-spot improvisations.

Over the past decades, considerable robotic research [4] has aimed to develop task-specific systems to automate trades in construction, such as floor leveling [5], spray painting [6], and bricklaying [7]. However, these robot-centric solutions often struggle with real-world integration due to the varied and flexible nature of construction work, which encompasses a wide array of tasks. These systems frequently lack the adaptability needed to handle every nuance of the job and thus fail to provide in-situ solutions comparable to human workers. Moreover, the practical challenges of operating, maintaining, and supplying these robotic systems on active construction sites, as opposed to controlled lab settings, remain significant hurdles.

Drawing from the lessons of past endeavors in construction robotics [8], our research adopts a human-centric approach, informed by the latest developments in deep reinforcement learning (RL) [9]. This approach diverges from the traditional aim of complete automation of manual construction processes. Instead, our study introduces an ecological integration of mobile robots into existing manual workflows, positioning them in assistive and supportive roles. By undertaking tasks that are physically demanding yet seemingly minor, such robots can emerge as mobile work companions, allowing human workers to concentrate on the more skilled and critical aspects of their work. This hybrid, transitive model of robotically supported work collaboration seeks to address enduring construction challenges, including reducing physical strain, mitigating workplace injuries, and enhancing workflow fluency.

Fig. 1: Site observation captures of the carpentry formwork activitity and construction site condition

Based on in-depth observations at an actual construction site (Fig. 1), we focused on carpentry formwork—a labor-intensive and prevalent construction activity—as a prototypical scenario to explore the envisioned robotic support. Informed by a qualitative study of the social and material specifics of this work scenario, we developed a prototype, termed ”work companion rover”, aimed at providing tangible support to a duo of carpentry formwork installation workers. This rover's support functions encompass autonomous delivery of tools and materials, weight-bearing capabilities, and companionship during work tasks. Following its development, the prototype underwent both qualitative and quantitative evaluations to assess its support capabilities, conducted both in lab settings and on an actual construction site with workers.

This paper's contributions are threefold: (1) introducing a human-centered "work companion rover" prototype, specifically designed to closely support carpentry workers in their existing, labor-intensive tasks, (2) developing a lightweight, modular, and expandable framework driven by RL-based social navigation methods that can foster safe and comfortable navigation of mobile robots in real-world construction environments, broadly construed, and (3) showcasing a practical and efficient pipeline for contextually aligning and improving generically pretrained RL models with context-

---

[1] Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA, USA. {yuningw, jwei3, dcardoso, hyaejino}@andrew.cmu.edu

979-8-3503-7769-9/24/$31.00 ©2024 IEEE

specific features. At a broader level, we aim to illustrate through these contributions a feasible, alternative pathway for integrating autonomous robots into construction and labor-intensive work. This approach strategically values human skill and expertise, while concurrently harnessing AI and robotics to improve work safety and workflow fluency.

II. BACKGROUND AND RELATED WORKS

A. *Assistive robots in various domains*

In robotics, decades of research and development have been dedicated to assistive robots and human-robot teaming across various domains [10], including healthcare [11], [12], public guiding and touring [13], disaster response [14], and space exploration [15]. These technologies have played a critical role in enhancing service quality and user experience. However, their application in labor-intensive sectors like on-site construction remains limited, despite construction being the largest global industry and one heavily reliant on manual labor in complex environments.

Our recent investigation at an active construction site revealed considerable potential for robotic assistance. Introducing such systems can reduce manual traversals, alleviate physical fatigue through sensitive support, and improve safety by reducing the risk of workplace injuries. During observations of eleven trades in the "rough carpentry stage"—a phase marked by concrete floor completions and intensive manual installations—we focused on carpentry formwork, a labor-intensive task involving a two-person team installing formwork panels to mold concrete walls. In addition to their primary tasks, workers spent significant time transporting tools and materials in small batches from a distant workbench or previous work area. We recorded approximately 19 such traversals in a single hour. This frequent activity represents an overlooked physical burden and a prime opportunity for the introduction of a "work companion rover"—a robot that could autonomously deliver tools and materials, and unobtrusively accompany workers by carrying heavy loads.

B. *Robot navigation and mobility on construction site*

Existing research in the field of robot navigation has shown considerable progress over the past decades [16]. However, the majority of the research, with a few notable exceptions, is based on contexts that differ from ours in terms of robotic functionality and application. For example, Logistic robots, which are increasingly common in factories and warehouses, are designed for efficient goal-directed deliveries in environments that are far more orderly and organized than the chaotic and confined spaces of indoor construction sites. A substantial amount of research in social robot navigation focuses on enabling robots to navigate among densely populated areas with social compliance [17], with applications such as robot delivery on urban streets [18] and university campuses [19]. In the construction domain, recent advancements in quadruped robots [20] demonstrate their potential to autonomously inspect, monitor, and document site progress, even across multiple floors. However, these capabilities are not tailored to directly provide human-centered, integrated support to heavy manual laborers within their current workflows.

Acknowledging these fundamental differences in context and functions, we identified two primary challenges specific to our research. Firstly, the robot must be capable of safely navigating through physically unstructured and complex construction environments, especially the floor area relevant to carpentry formwork. Secondly, to effectively provide support to carpentry workers, the robot should be able to navigate comfortably and closely around workers, taking into account their specific work activities and features. Tackling these challenges demands a nuanced understanding of both the physical and social aspects of robot navigation within the unique context of on-site construction work.

C. *Reinforcement learning for social robot navigation*

Recent advances in social robot navigation [21] have addressed many challenges related to the social aspects of navigation. Notably, reinforcement learning has emerged as a promising method for adapting to the dynamic movements of neighboring agents. Various deep neural network (DNN)-based algorithms have been developed to tackle these social dynamics. For example, early studies [22] used Long Short-Term Memory (LSTM) networks to capture latent encodings of neighboring agents' movements. Later research [23] extended this to model both human-human and human-robot interactions, enriching the learning process. More recent work [24] employed decentralized structural Recurrent Neural Networks (RNNs) to explicitly model spatial and temporal relationships within crowds. Additionally, other studies [25] focused on first-person perspectives, benchmarking the impact of this limited viewpoint on navigation performance. Despite these advancements, gaps remain in applying these methods effectively to our specific research context. Much of this work assumes open spaces with few physical obstacles and is situated in crowd-based settings where agents are primarily modeled as goal-seeking pedestrians.

The core challenge of our research is to meaningfully leverage existing RL-based social robot navigation approaches while adapting them to the intricacies of construction work. This involves not only designing a framework to account for complex site surroundings but also ensuring that social dynamics of human-robot interaction remain effective in more challenging and task-driven work settings. Unlike goal-seeking pedestrians, construction workers exhibit activity patterns that are far more varied and task-specific. Successfully integrating these aspects—navigating a physically complex environment while adapting to worker-specific interactions and needs—will be crucial to our research.

III. PROBLEM STATEMENT

A. *Robotic support scenarios and functions*

To realize and validate the envisioned robotic support, we identified three key functions across two primary scenarios. The first scenario involves tool/hardware delivery between a distant workbench and the workers' current work zone (left in Fig. 2), while the second focuses on load-bearing and accompanying workers as they move between adjacent zones (right in Fig. 2). The essential robot functions are *send*, *summon*, and *accompany*. The *send* function dispatches the robot to a central workbench (approx. 20m away) to retrieve tools and materials. The *summon* function remotely recalls the robot once the load is ready. The *accompany* function allows the robot to assist workers by carrying tools and materials as they transition to new work zones. These functions prioritize navigation comfort, ensuring the robot does not obstruct or interrupt ongoing work, and are designed to be operated via simple, intuitive commands like a single key press on a controller.

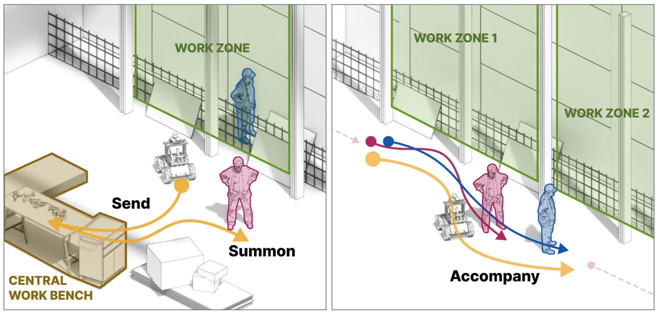

Fig. 2: The robotic support scenarios and functions

B. MDP formulation

The essence of the above robot scenarios and functions is modeled as a 2D navigation task where a robot moves towards an anticipated goal position while encountering and interacting with n workers in a cluttered physical layout L. Following a similar formulation in [23], [24], we model the worker-robot interaction scenarios as a Markov Decision Process (MDP) of $\langle \mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}, \gamma, \mathcal{S}_0 \rangle$. The observable state for both worker and robot includes position $\mathbf{p} = [p_x, p_y]$, $\mathbf{v} = [v_x, v_y]$, and safety radius r. For the robot, we assume it is aware of its preferred speed v_{pref} and goal position $\mathbf{p}_g = [p_{gx}, p_{gy}]$, and that $\mathbf{v}_t = \mathbf{a}_t$. At time t, the joint state for the robot encompasses both its own full state $\mathbf{s}_t^{\text{full}}$ and the observable states of all neighboring workers $\mathbf{w}_t^{\text{obs}}$, $\mathbf{s}_t^{\text{joint}} = [\mathbf{s}_t^{\text{full}}, \mathbf{w}_t^{\text{obs}}]$. The MDP can be solved through reinforcement learning to maximize the expected return, $R_t = \mathbb{E}\left[\Sigma_{k=t}^{T} \gamma^{k-t} r_k\right]$. Our goal is to optimize this objective while taking contextual information of cluttered layout L and workers' unique activity patterns into consideration.

IV. APPROACHES

In response to the challenges described, our research adopts a research-by-prototyping approach. Drawing from site observations and worker interviews, we developed a mid-sized companion rover prototype based on a generic UGV platform, designed for close-range operation around workers. This prototype was used for real-world experimentation of the envisioned support. The prototyping process began with the creation of a modular system framework (Fig. 3), built for flexibility and adaptability to the specific demands of a construction work environment. This framework is based on the ROS navigation stack with context-specific customizations and incrementally integrates modules to address (1) the complexities of the construction site, (2) the perception of carpentry workers' activities, and (3) the contextual adaptation of RL-based social navigation methods. The framework includes five key modules: (1) hardware system design, (2) site mapping and robot state estimation, (3) worker detection and tracking, (4) hierarchical motion planning, and (5) contextual fine-tuning of RL. Each module is detailed below with a focus on addressing the unique challenges of this domain.

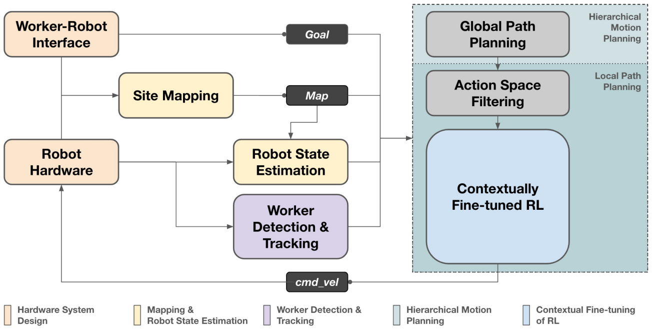

Fig. 3: The modular framework design

A. Hardware system design

The challenging conditions of a construction site, combined with the demands of tool and material transport, required careful consideration in the design of our robot's hardware (Fig. 4). During the carpentry formwork stage, the construction floor, though mostly flat, is often cluttered with dust, puddles, cables, and low-lying obstacles such as hand tools and material piles. To navigate effectively over such rough terrain, the robot chassis needs to be both robust and agile. Rather than building a chassis from scratch, we determined that many existing mid-sized UGV bases could meet these navigation needs, allowing us to focus on more context-specific challenges rather than chassis development. With this in mind, we evaluated three different robot bases: a Fetch robot, a custom-built omnidirectional robot, and a Clearpath Husky robot.

Our trials showed that the Fetch robot and the omnidirectional robot, both hindered by smaller wheel sizes, struggled to maneuver over cables, potholes, and coarse sawdust. In contrast, the Clearpath Husky, designed for rugged outdoor terrains, proved to be the superior choice for our application. Equipped with large, durable wheels and differential control motors, the Husky robot navigated uneven surfaces and obstacles with ease. Its low center of gravity further enhanced stability, critical when carrying tools and materials. Additionally, the inclusion of two T-slot rails on the Husky provided flexibility for installing additional equipment, making it well-suited to the demands of our context.

The sensor package atop the robot base is a multi-modal assembly designed for navigating the geometric complexities of a construction site and perceiving worker activities. A 3D LiDAR (Velodyne VLP-16) is positioned on top of the robot

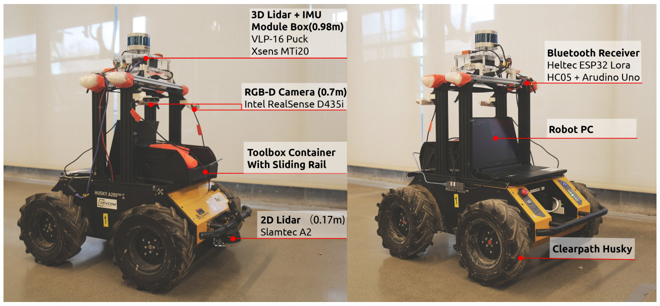

Fig. 4: Hardware design of the robot

frame for maximal point cloud data collection, coupled with an Xsens IMU. A 2D LiDAR (RPLiDAR A2), placed at the robot's front bottom, compensates for the 3D LiDAR's blind zone, ensuring the detection of objects as low as 8cm (e.g., hand tools, materials, and large debris) in the robot's path. Additionally, two Intel RealSense D435i RGB-D cameras are installed on the sides for effective worker detection and tracking.

The robot integrates an onboard computer with an NVIDIA RTX 3070 Ti GPU for real-time inference of multiple DNNs. It also features a 35cm x 45cm retractable container on the deck, sized based on common carpentry tools and materials. After sensor addition, the maximum payload of the robot is 35 kg on concrete surfaces, accommodating up to 2-3 bags of metal hardware and 2-3 types of hand tools like rotary drills, pliers, and hammers. Additionally, a Bluetooth keypad controller customized for carpentry workers enables the robot to operate conveniently even with work gloves, enhancing worker-robot interaction in noisy environments.

B. *Mapping and localization on unstructured site*

The robot's effectiveness in supportive navigation is fundamentally anchored in quality mapping and robot state estimation within the geometrically sophisticated construction work environment. Objects on the site range from very tall (e.g., ladders, steel columns, frames, material stacks, boxes) to very low (e.g., hand tools, material piles, waste clusters, randomly placed extension sockets), and are inevitably distributed in an unpredictable manner. This setting diverges significantly from more structured environments like warehouses, manufacturing factories, or open public spaces.

Conventional 2D mapping tools, such as `gmapping`, are insufficient for capturing the complex topography of a construction site (Fig. 5), such as height-varying objects (e.g., open ladders, tool tripods) and randomly extruding hazards (e.g., horizontally stacked steel beams and pipes). Meanwhile, 3D voxel-based or point cloud-based representations are too computationally inefficient for downstream navigation tasks. To address this, we employ a recent LiDAR Odometry and Mapping (LOAM) [26], [27] method to generate a comprehensive 3D point cloud of the site by manually navigating the robot around the intended operational area with a joystick for 1-2 loops. The resulting point cloud is then projected and processed into a 2D grid map, effectively balancing the accuracy of 3D mapping with the computational efficiency of 2D methods, crucial for capturing environmental features in our context.

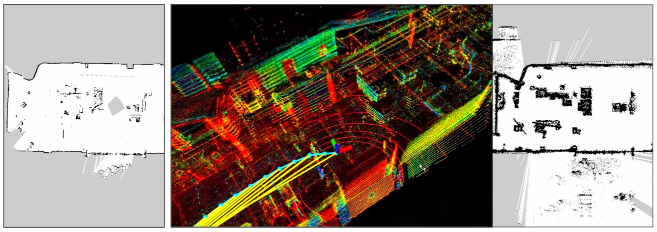

Fig. 5: Standard 2D grid map (left) compared to our condensed map (right)

The cluttered and dynamic nature of construction sites also poses a challenge for the robot's state estimation, i.e., accurately determining its position and velocity with the processed map. Conventional scan matching methods like AMCL are prone to errors and drifts due to the non-stationary nature of construction objects. These objects, such as material piles, are temporary and change over time. To overcome this besides constant map updates, we integrated multiple odometry and corrective information sources into the localization process through an Extended Kalman Filter (EKF), including IMU data, odometry from LOAM, the robot's wheel encoders, and AMCL.

C. *Worker detection and tracking*

To ensure comfortable navigation around carpentry workers, it is imperative that the robot promptly recognizes and tracks nearby workers' movements (i.e., position, velocity, and radius) against a complex backdrop, especially within its immediate action space. Drawing insights from [28], we break down this task through detection and tracking.

Our qualitative physical experiments revealed that detection methods relying solely on 3D LiDAR data, although effective in contexts like autonomous driving and crowd detection, falter in our context. The failure is caused by the unstructured and unpredictable nature of construction objects. In many cases, the scan appearance of randomly stacked materials can mimic that of a person, leading to a high false positive rate. As an alternative, we found that vision-based methods offer greater reliability and more computational efficiency, particularly when a worker's face is within the camera's field of view (FoV). Consequently, we adopted YOLOv7 [29] for worker detection, and utilized depth information to estimate the workers' position, velocity, and radius. To overcome a single RGB-D camera's limited FoV, we installed two cameras and amalgamated their detection results by clustering [30]. Considering that the robot's low height can hinder the visibility of a worker's face when they are close to the robot, we tilted the RGB-D cameras upwards by 15 degrees to physically maintain a better vantage point. As a last step, the raw detection results were processed with Multi-Object Tracking (MOT) algorithms [31] to reduce noise and ensure worker tracking continuity, especially given frequent visual obstructions by site objects and structures.

D. Hierarchical motion planning

The robot's motion planning utilizes a hierarchical approach, integrating both search-based methods and RL-based techniques for navigation in cluttered, populated environments. For global planning, the robot employs A* [32] to generate a path based on the latest condensed map. At the local path planning stage, a lightweight search-based method, such as the Dynamic Window Approach (DWA) [33], acts as a foundational safety layer by efficiently filtering out actions that exceed the robot's dynamic constraints or pose collision risks (Eq. 1). Rather than relying solely on the efficiency-driven scoring mechanisms typical of search-based methods, a trained RL value network is used. This network, optimized for social comfort and compliance, samples and selects the most appropriate final action from the refined action space $\mathcal{A}_f$ (Eq. 2).

$$\mathcal{A}_f = \mathcal{A} \cap \mathcal{A}_{\text{admissible}} \cap \mathcal{A}_{\text{dynamic}} \tag{1}$$

$$\mathbf{a}_t \leftarrow \text{argmax}_{a_t \in \mathcal{A}_f} R\left(\mathbf{s}_t^{\text{joint}}, \mathbf{a}_t\right) + \gamma^{v_{\text{pref}} \cdot \Delta t} V_{t+\Delta t}^{\text{joint}} \tag{2}$$

This hierarchical approach offers three key advantages: generalizability, efficiency, and explainability. For generalizability, the approach enables flexible comparisons between different search-based and RL methods in real-world experiments with minimal engineering adjustments. In terms of efficiency, the hierarchical framework naturally falls back to a search-based method when no workers are detected, making it highly efficient in deployment for low-traffic environments and energy conservation. To consolidate this efficiency and ensure smooth integration between layers, we incorporated measures such as fine-tuning the searching sparsity of DWA and accelerating parallel value network inference for action selection. Finally, for explainability, compared to an end-to-end reinforcement learning method, this layered approach allows for clear diagnostics of issues leading to collisions during prototyping. When the robot collides with either humans or objects, the hierarchical framework enables pinpointing exactly which layer of the system failed. As we further consolidate the prototype, reinforcement learning approaches that account for both site objects and worker movements represent a promising future direction.

E. Contextual training and fine-tuning of RL algorithm

Through preliminary on-site observations, we identified distinct worker behavior attributes compared to those typically assumed in crowd-based social navigation research. (1) Worker groups are small (fewer than 5) but denser in compact spaces. (2) Their movements are highly task-related, resulting in unique behavioral patterns and greater movement diversity and uncertainty compared to pedestrians. (3) Engrossed in tasks, workers have limited attention to spare for the robot. These attributes underscore the key safety challenges in applying existing methods to this context.

Acknowledging these differences, we leverage existing RL-based social navigation research as a foundation while incorporating new insights inspired by the context. In a simulation environment [23], we first conducted pre-training using generic setups commonly found in social navigation research, such as circle and square crossing [34], [23]. We evaluated a pool of RL-based social navigation methods, including CADRL [34], GA3C-CADRL [22], LSTM-RL [23], SARL [23], and others, applying them to the robot for an initial round of selection. In this preliminary selection, SARL qualitatively outperformed other models in delivering smoother navigation and was chosen for subsequent fine-tuning and alignment.

Drawing insights from firsthand observations of workers' daily activities on-site, we conducted an additional phase of contextual fine-tuning on the selected pretrained model (Fig. 6). This fine-tuning involved designing a series of experimental setups that incorporated agent behavioral attributes and diversity not typically included in previous social navigation research but critical in workers' daily practices. For instance, some customized patterns include: (1) "stop-and-go", where intentional, intermittent pauses mimic typical worker behavior of walking along a work wall and taking brief pauses for installation tasks, (2) "back-and-forth", a common pattern where workers repeatedly move short distances between a work spot and a temporary stash, and (3) "in-place walkabouts", where workers move within a confined area for focused installation work, displaying limited but non-static movement. Along with these patterns, generated through rule-based simulation, we also introduced variables such as fluctuating agent safety radii, limited fields of view, denser agent initialization, and intentional close encounters.

While these customizations introduced greater uncertainty and diversity in training, it was equally important to fine-tune the pretrained model without causing drastic policy shifts. Maintaining the navigation capabilities acquired during generic training is crucial to avoid overfitting and to ensure the model's stability for Sim2Real deployment. To achieve this, we innovatively adapted the reward function r_{ori} from [23] by introducing a shift penalization mechanism. Actions $\{\mathbf{a}_1, \cdots, \mathbf{a}_m\}$ sampled from the refined action space $\mathcal{A}_f$ use their estimated values as a score for "preferability". Using the Plackett-Luce model [35] (Eq. 3), we convert value rankings into a probabilistic distribution (Eq. 4), penalizing policy shifts through KL-divergence to ensure gradual and stable model adaptation (Eq. 5).

$$P(\mathbf{a}_i|V) = \frac{V(\mathbf{a}_i)}{\Sigma_{j=1}^{m} V(\mathbf{a}_j)} \tag{3}$$

$$D_{\text{KL}}(\pi_{\text{old}} \| \pi_{\text{new}}) = \Sigma_{i=1}^{m} P_{\text{old}}(\mathbf{a}_i) \log\left(\frac{P_{\text{old}}(\mathbf{a}_i)}{P_{\text{new}}(\mathbf{a}_i)}\right) \tag{4}$$

$$r = r_{\text{ori}} - \lambda D_{\text{KL}} \tag{5}$$

V. EXPERIMENTS AND RESULTS

To physically examine the robot's performance in offering tangible support, we conducted the real-world evaluation and demonstration of the robot both qualitatively and quantitatively (1) on an actual construction site with workers

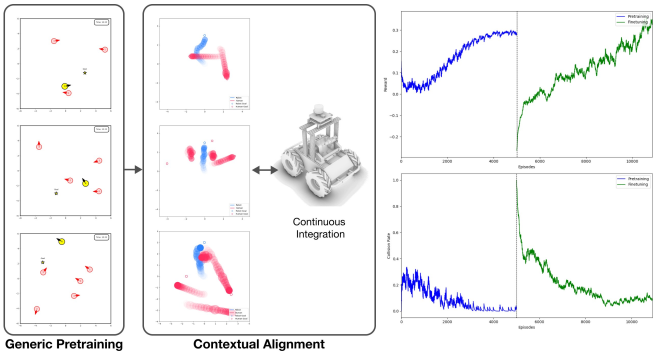

Fig. 6: Contextual curriculum alignment

and (2) in the lab sandbox. In the on-site demonstration and evaluation phase, we transported the robot to an active construction site and tested the robot's functions through direct interaction with the workers. In the lab evaluation, we drew critical insights from the on-site demonstration and conducted a detailed quantitative analysis regarding several metrics.

A. *On-site demonstration and evaluation*

In our on-site demonstration and evaluation, we set up a 45m x 35m carpentry formwork demonstration area on an active construction floor (Fig. 7). This setup included a target formwork wall and several sequential work zones along it. About 15m from the wall and work zones, we designated a central workbench storing tools and metal hardware. Three construction workers participated in the evaluation, performing typical formwork installation tasks and interacting with the robot in three scripted scenarios. These scenarios, derived from our Problem Statement, involved: (1) directing the robot from the current work zone to the central workbench for tools/hardware retrieval, (2) remotely summoning the robot from the workbench to their work zone after loading, and (3) keeping the robot nearby to assist in transporting tools and hardware across installation zones. Each scenario was executed twice. The research team minimally intervened during these demonstrations, focusing primarily on documentation through cameras. The workers operated the robot using the custom-built keypad controller. Additionally, we introduced an extra test to challenge the robot's navigational capabilities in sparsely mapped and unfamiliar areas.

The robot effectively met its support and companion roles in the initial three scenarios. Our team qualitatively noted occasional navigation pauses, particularly when the robot encountered passageways near its threshold width amidst construction objects. Regarding social comfort, there were no significant disruptions observed in the workers' activities within the active zones, though we deferred to the workers for their perspectives on this aspect. The robot typically adjusted its path, maintaining a distance of one to two steps from the workers during active navigation. In cases where a team member deliberately and suddenly obstructed its path, the robot promptly halted and executed a brief detour. The robot, however, terminated early in the extra test case due to a lack of detailed mapping, and an unsuccessfully detected extruded pipes, converging to a localization error.

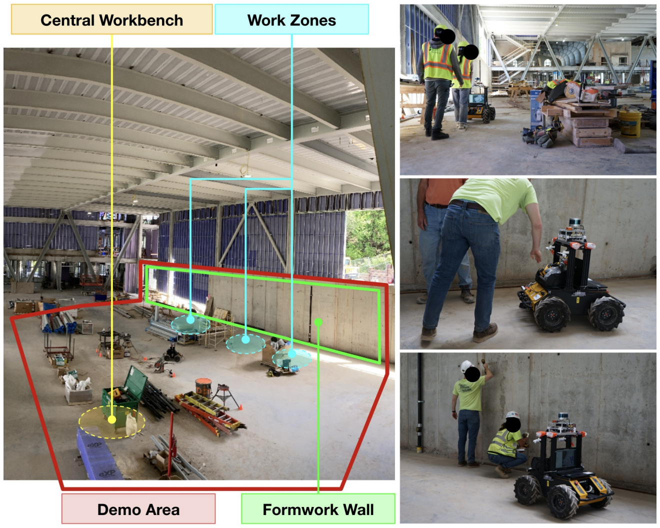

Fig. 7: Demo and evaluation on an actual construction site

After the demonstration, we conducted a critical qualitative evaluation through a group interview with the participating workers. This interview covered three technical aspects of the robot, each linked to the context-specific challenges identified earlier.

- **Robot navigation safety and comfort**: The questions posed included inquiries about the robot's navigation speed, the safety of working around it, and comfort levels, including any obstructions caused by its presence. Workers provided concise yet affirmative responses, generally agreeing that it was safe and comfortable to work alongside the robot. They did not perceive the robot as abrupt or intimidating and found it a seamless addition to their workflow.
- **Robot hardware design and usability**: Questions here revolved around the robot's hardware and system design, operational ease, and the intuitiveness of the controller. Workers acknowledged the robot's user-friendly interaction design and non-threatening appearance. They particularly appreciated our efforts to cover sharp edges with bubble foam for safety. However, there was a desire for enhanced physical capabilities, such as increased load capacity and the ability to haul carts.
- **Overall concept of robotic support in current workflow**: We inquired about the usefulness of the robot in carpentry formwork and other potential applications in construction. Workers validated the robot's current utility in reducing on-site travel and suggested its applicability in tasks like hardware distribution during interior decoration phases. They also expressed interest in future enhancements, such as autonomous tool/material loading features with robotic arms, suggesting potential for expanded manipulation functionalities.

These insights not only affirmed our initial proposal but also opened up intriguing avenues for future research and development in the realm of supportive and assistive work companion robots in heavy manual construction work.

B. *Quantitative lab evaluation*

In our laboratory, we carried out quantitative analyses and ablation studies to further validate the robot's support capabilities, particularly assessing the necessity of employing and fine-tuning RL in our envisioned support scenario (Fig. 8). This included a performance comparison between the robot operating with and without the RL-based social navigation layer, contrasting the RL-driven hierarchical motion planning framework with basic efficiency-driven collision avoidance (Fig. 9).

A sandbox environment was established, maximally mirroring the on-site conditions and challenges. Evaluation metrics focused on the rate of comfortable encounters and delivery, defined by criteria such as worker proximity, movement abruptness, navigation delays, unwanted surprises, and work obstructions. Additional metrics like collision rate, success rate, freeze rate, dangerous movement rate, and travel time were also documented for a holistic performance assessment.

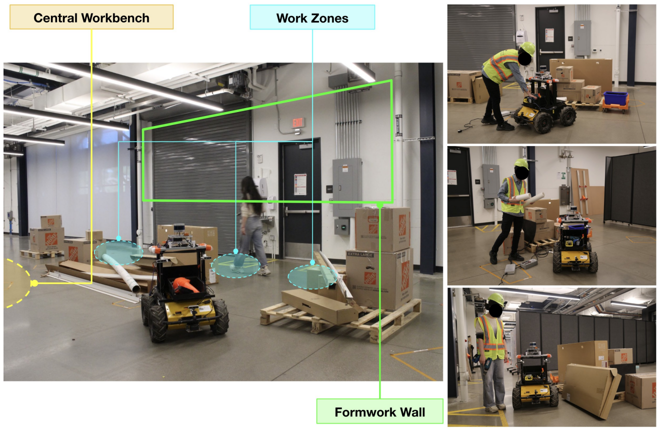

Fig. 8: Lab sandbox environment evaluation

We designed two scenarios with 10 distinct test cases each, reflecting diverse worker activities and robot objectives. Two team members with context knowledge and safety training simulated worker behaviors—working beside walls, standing still, and moving around the site. Each case was executed six times, alternating among contextual RL, vanilla RL, and non-RL runs, in a randomized, blind test setup to minimize subjective bias in evaluating the robot's interactions. The robot's performance was carefully assessed against aforementioned criteria. Particularly, the first scenario assessed the robot's ability to navigate between the central workbench and work zones around workers amidst obstacles, while the second evaluated its capacity to accompany workers in constrained spaces between work zones.

Our empirical findings (Table I) underscore the effectiveness and necessity of the RL-based social navigation layer for enhancing the robot's support capabilities in cluttered, worker-centric environments. In scenarios not enacting RL, the comfortable encounter and delivery rates markedly decreased by approximately 0.4 and 0.15 respectively, While RL implementation did not significantly reduce major/minor collision rates (Table II), it substantially decreased obstructions to workers, robot freezes and tangentially dangerous movements, resulting in smoother and more socially aware and comfortable support around workers, which validated the initial design intention.

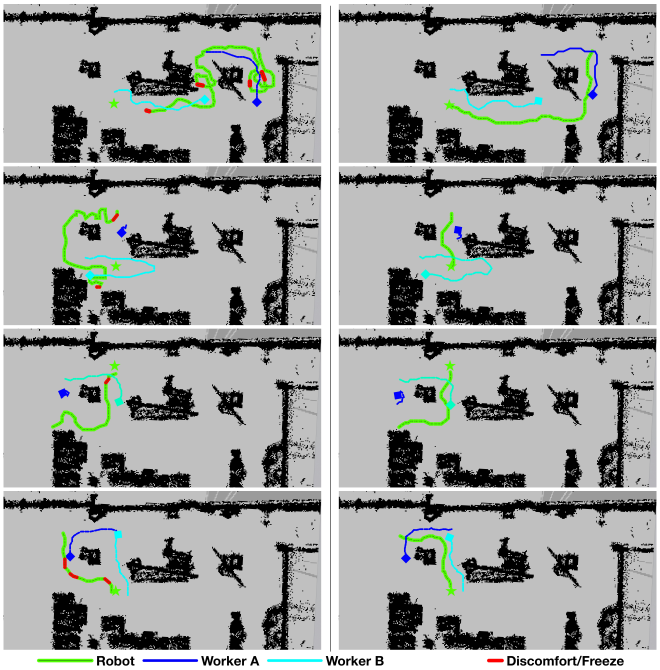

Fig. 9: Trajectory comparison between RL disabled (left) and enabled (right)

	Comf. Delivery	Comf. Encounter
DWA	0.83	0.57
DWA+SARL (Vanilla)	0.95	0.87
DWA+SARL (Contextual)	**0.98**	**0.92**

TABLE I: Quantitative results in lab evaluation regarding comfort-related metrics.

	Success	Collision	Freeze	Time
DWA	0.83	0.05	0.95	29.40
DWA+SARL (Vanilla)	1.00	0.07	**0.23**	25.59
DWA+SARL (Contextual)	1.00	0.05	**0.10**	25.89

TABLE II: Quantitative results in lab evaluation. "Success": successful delivery rate. "Collision": major/minor collision rate. "Freeze": robot freeze rate. "Time": navigation time (s)

Further analysis reveals that our contextualized RL model, compared to the vanilla RL model trained only in generic settings, shows an improvement in the comfortable encounter rate and a decrease in the robot freeze rate, resulting in less obstruction or interruption when offering support. This enhancement is attributed to the additional alignment tailored to work-specific scenarios and cases, which better equips the robot to handle unique patterns and attributes in our context.

VI. CONCLUSION AND FUTURE WORK

This study marks a considerate departure from robot-centric automation in the heavily manual construction setting, introducing a human-centric "work companion rover" designed to assist existing human effort. Leveraging advancements in reinforcement learning, the rover is deeply contextualized in the material and social aspects of an actual construction work, in the hope of offering tangible and meaningful help to heavily manual carpentry workers. Our empirical results underscore the robot's effectiveness in real-world scenarios, offering a novel assistive prototype while highlighting a promising avenue for more diverse future research in similarly labor-intensive work contexts. Besides the current robotic support function and scenario, the developed robotic framework is possible to serve as an expandable navigation base for other supportive scenarios, approaches, and applications, calling for a future of labor ecologically combining human skill and expertise with advances in AI and robotics.

REFERENCES

[1] M. E. Shehata and K. M. El-Gohary, "Towards improving construction labor productivity and projects' performance," *Alexandria Engineering Journal*, vol. 50, no. 4, pp. 321–330, Dec. 2011.

[2] R. K. Fagbenro, R. Y. Sunindijo, C. Illankoon, and S. Frimpong, "Influence of Prefabricated Construction on the Mental Health of Workers: Systematic Review," *European Journal of Investigation in Health, Psychology and Education*, vol. 13, no. 2, pp. 345–363, Feb. 2023.

[3] P. Wu, J. Wang, and X. Wang, "A critical review of the use of 3-D printing in the construction industry," *Automation in Construction*, vol. 68, pp. 21–31, Aug. 2016.

[4] M. Gharbia, A. Chang-Richards, Y. Lu, R. Y. Zhong, and H. Li, "Robotic technologies for on-site building construction: A systematic review," *Journal of Building Engineering*, vol. 32, p. 101584, Nov. 2020.

[5] T. Liu, H. Zhou, Y. Du, J. Zhang, J. Zhao, and Y. Li, "A Brief Review on Robotic Floor-Tiling," in *IECON 2018 - 44th Annual Conference of the IEEE Industrial Electronics Society*. Washington, DC: IEEE, Oct. 2018, pp. 5583–5588.

[6] E. Asadi, B. Li, and I.-M. Chen, "Pictobot: A Cooperative Painting Robot for Interior Finishing of Industrial Developments," *IEEE Robotics & Automation Magazine*, vol. 25, no. 2, pp. 82–94, June 2018.

[7] Z. Fang, Y. Wu, A. Hassonjee, A. Bidgoli, and D. Cardoso Llach, "Towards an Architectural Framework for Distributed, Robotically Assisted Construction: Using Reinforcement Learning to Support Scalable Multi-Drone Construction in Dynamic Environments." in *Proceedings of the 40th Annual Conference of the Association of Computer Aided Design in Architecture (ACADIA)*, vol. I: Technical Papers. Online and Global: CUMINCAD, Oct. 2020, pp. 320–329.

[8] T. Bock, "The future of construction automation: Technological disruption and the upcoming ubiquity of robotics," *Automation in Construction*, vol. 59, pp. 113–121, Nov. 2015.

[9] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*. MIT press, 2018.

[10] L. Mingyue Ma, T. Fong, M. J. Micire, Y. K. Kim, and K. Feigh, "Human-Robot Teaming: Concepts and Components for Design," in *Field and Service Robotics*, M. Hutter and R. Siegwart, Eds. Cham: Springer International Publishing, 2018, pp. 649–663.

[11] L. V. Calderita, L. J. Manso, P. Bustos, C. Suárez-Mejías, F. Fernández, and A. Bandera, "THERAPIST: Towards an autonomous socially interactive robot for motor and neurorehabilitation therapies for children," *JMIR rehabilitation and assistive technologies*, vol. 1, no. 1, p. e1, Oct. 2014.

[12] J. Fasola and M. J. Matarić, "A socially assistive robot exercise coach for the elderly," *Journal of Human-Robot Interaction*, vol. 2, no. 2, pp. 3–32, June 2013.

[13] T. Iio, S. Satake, T. Kanda, K. Hayashi, F. Ferreri, and N. Hagita, "Human-Like Guide Robot that Proactively Explains Exhibits," *International Journal of Social Robotics*, vol. 12, no. 2, pp. 549–566, May 2020.

[14] G.-J. M. Kruijff, M. Janíček, S. Keshavdas, B. Larochelle, H. Zender, N. J. Smets, T. Mioch, M. A. Neerincx, J. V. Diggelen, F. Colas, *et al.*, "Experience in system design for human-robot teaming in urban search and rescue," in *Field and Service Robotics: Results of the 8th International Conference*. Springer, 2014, pp. 111–125.

[15] S. Hopko, J. Wang, and R. Mehta, "Human factors considerations and metrics in shared space human-robot collaboration: A systematic review," *Frontiers in Robotics and AI*, vol. 9, p. 799522, 2022.

[16] A. Pandey, S. Pandey, and DR. Parhi, "Mobile robot navigation and obstacle avoidance techniques: A review," *Int Rob Auto J*, vol. 2, no. 3, p. 00022, 2017.

[17] C. Mavrogiannis, F. Baldini, A. Wang, D. Zhao, P. Trautman, A. Steinfeld, and J. Oh, "Core Challenges of Social Robot Navigation: A Survey," Mar. 2021.

[18] C. Chen, E. Demir, Y. Huang, and R. Qiu, "The adoption of self-driving delivery robots in last mile logistics," *Transportation Research Part E: Logistics and Transportation Review*, vol. 146, p. 102214, Feb. 2021.

[19] J. Williamsson, "Business model design for campus-based autonomous deliveries – A Swedish case study," *Research in Transportation Business & Management*, vol. 43, p. 100758, June 2022.

[20] EM. Wetzel, J. Liu, T. Leathem, and A. Sattineni, "The use of boston dynamics SPOT in support of LiDAR scanning on active construction sites," in *ISARC. Proceedings of the International Symposium on Automation and Robotics in Construction*, vol. 39. IAARC Publications, 2022, pp. 86–92.

[21] P. T. Singamaneni, P. Bachiller-Burgos, L. J. Manso, A. Garrell, A. Sanfeliu, A. Spalanzani, and R. Alami, "A survey on socially aware robot navigation: Taxonomy and future challenges," *The International Journal of Robotics Research*, p. 02783649241230562, Feb. 2024.

[22] M. Everett, Y. F. Chen, and J. P. How, "Motion planning among dynamic, decision-making agents with deep reinforcement learning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Madrid, Spain, Sept. 2018.

[23] 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)*. IEEE, 2019, pp. 6015–6022.

[24] S. Liu, P. Chang, W. Liang, N. Chakraborty, and K. Driggs-Campbell, "Decentralized Structural-RNN for Robot Crowd Navigation with Deep Reinforcement Learning," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. Xi'an, China: IEEE, May 2021, pp. 3517–3524.

[25] B. Stoler, M. Jana, S. Hwang, and J. Oh, "T2FPV: Dataset and method for correcting first-person view errors in pedestrian trajectory prediction," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 4037–4044.

[26] J. Zhang and S. Singh, "LOAM: Lidar Odometry and Mapping in Real-time," in *Robotics: Science and Systems X*. Robotics: Science and Systems Foundation, July 2014.

[27] T. Shan, B. Englot, D. Meyers, W. Wang, C. Ratti, and D. Rus, "LIO-SAM: Tightly-coupled Lidar Inertial Odometry via Smoothing and Mapping," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Oct. 2020, pp. 5135–5142.

[28] M. F. Everett, "Robot designed for socially acceptable navigation," Ph.D. dissertation, Massachusetts Institute of Technology, 2017.

[29] C.-Y. Wang, A. Bochkovskiy, and H.-Y. M. Liao, "YOLOv7: Trainable bag-of-freebies sets new state-of-the-art for real-time object detectors," *arXiv preprint arXiv:2207.02696*, 2022.

[30] E. Schubert, J. Sander, M. Ester, H. P. Kriegel, and X. Xu, "DBSCAN revisited, revisited: Why and how you should (still) use DBSCAN," *ACM Transactions on Database Systems (TODS)*, vol. 42, no. 3, pp. 1–21, 2017.

[31] N. Wojke, A. Bewley, and D. Paulus, "Simple online and realtime tracking with a deep association metric," in *2017 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2017, pp. 3645–3649.

[32] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.

[33] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.

[34] Y. F. Chen, M. Everett, M. Liu, and J. P. How, "Socially aware motion planning with deep reinforcement learning," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sept. 2017, pp. 1343–1350.

[35] R. L. Plackett, "The analysis of permutations," *Journal of the Royal Statistical Society Series C: Applied Statistics*, vol. 24, no. 2, pp. 193–202, 1975.