Ph.D. Thesis Proposal

Towards Humane Automation: An RL-Driven Robotic Framework for Supporting On-Site Construction Workers

Yuning Wu
Ph.D. Candidate in Computational Design
Carnegie Mellon University

Dissertation Committee

Dr. Daniel Cardoso Llach
Advisor
Associate Professor, Computational Design Track Chair
School of Architecture
Carnegie Mellon University

Dr. Jean Oh
Co-advisor
Associate Research Professor
Robotics Institute
Carnegie Mellon University

Dr. Jieliang (Rodger) Luo
External Committee Member
Principal AI Research Scientist
Autodesk Research, AI Lab

Carnegie Mellon University

Abstract

How might robots humanely support people in their existing labor-intensive work? Contextualized in a specific construction scenario, this research addresses this question through the development of a robotic framework for adaptively supporting construction workers. Employing a “research by prototyping” approach informed by insights drawn from qualitative studies of an actual construction site, the framework mobilizes state-of-the-art Reinforcement Learning (RL) algorithms and includes five main technical elements, (1) robot hardware prototyping, (2) unstructured site perception and robot state estimation, (3) worker detection and tracking, (4) hierarchical motion planning, and (5) contextualized RL training and deployment. Realized through an actual physical robot prototype, a “work companion rover” for carpentry workers, the framework is examined both qualitatively and quantitatively in lab settings and on an actual construction site. Instead of purporting to replace human workers, this research shows how advanced computational and robotics techniques might be designed to adaptively support them. It advances the state-of-the-art in computational design by innovatively bringing together AI and robotics techniques to actual construction work contexts rather than idealized lab settings. In addition, it also contributes to the fields of robotics and AI by demonstrating a human-centered use of RL for supporting heavy manual labor and showcasing a technology design process deeply informed by the social and material specificity of construction work.

Introduction

“Will robots replace humans?” (Fig. 1) This is the question I get asked the most on my morning Uber to the Mill-19 lab, once people hear my research domain. Beneath the frequently asked question lies the common eagerness for an answer about the future of robots. Instead of picturing a vision where nimble robots are designed to fully replace human labor, I wish to imagine the possibility of a different route, which I intend to describe as “Humane Automation”. By humane, it means more than just a friendly stance toward existing human efforts. It particularly envisions human-robot collaborative work capabilities in complex, labor-intensive scenarios, with on-site construction work being a highly representative context for the investigation.

Figure 1. “Hey ChatGPT, finish this building…” appeared on a building cover in Antwerp, Belgium. The irony implies the huge hidden opportunity underlying the seemingly omnificent AI model trending recently. ChatGPT has demonstrated strong capabilities in understanding human knowledge and answering difficult conversational questions. Yet it is hand-tied when dealing with physical activities in the real world. Human skills are still irreplaceable in trades such as finishing a building. (Source: Impact)

Construction is one of the world’s largest industries, yet it is one of the most labor-intensive, arduous, and least automated. Unlike other manufacturing contexts (e.g., airplanes, cars, and microchips) that can be neatly housed in organized factories and produced by pre-programmed machines in standardized pipelines, most construction activities must happen on-site in one way or another, with skilled human workers tending to highly bespoke needs in messy, complicated work environments. The canonical way of automation, which turns an atomic manual procedure into a single-purpose, task-specific autonomous system, faces challenges in the real world because of the myriad of cumbersome, improvisational, composite tasks occurring on-site. Instead of proposing yet another instance of such robotic systems, a fundamental rethinking of the roadmap toward automation is at stake. Drawing critical insights from recent advances in AI and robotics, my research aims to target the commonly assumed role of autonomous robots in on-site construction work. Namely, rather than developing robotic systems as human surrogates, can robots be introduced into existing labor activities such that they can accompany and support human workers while providing handy assistance? How can recent advances in AI and robotics be mobilized to offer the described support in adaptive and comfortable ways? Contextualized in a specific on-site construction labor activity, my research investigates this humane form of automation through hands-on research and development, which I summarize as a **Reinforcement Learning (RL) driven robotic framework for supporting on-site construction workers**. This framework is developed in tight coupling with my role as a member of the Rethinking Automation in Construction research group (ReAC) at the Computational Design Laboratory at Carnegie Mellon University SoA, where I led technical efforts towards what the group termed "robotically supported cooperative construction". The framework is realized and validated through an actual physical robot prototype, a "**work companion rover**", and is examined through both in-lab and on-site efforts qualitatively and quantitatively regarding its support capabilities.

Importantly, the research wishes to explore a less visited territory by former computational design researchers. When approaching the problem of how we build and make the physical space, previous research mostly tackles from a material-centered [1], design-centered [2], or machine-/robot-centered [3], [4] perspective. The substantial, longstanding presence of human labor, as well as the inextinguishable real-world messiness of work scenarios, are often simplified, idealized, or even "silenced" in research labs. This research, on the contrary, aims to face these hidden yet critical challenges head on, by proposing a human-centered technical architecture mobilizing state-of-the-art computational methods that can adapt into, rather than ideally reformulate, the existing workflows and contexts.

Notably, I must also acknowledge the rich body of existing and ongoing research in areas such as human-centered AI (HCAI) [5], [6], human-centered automation [7], collaborative automation [8], and socially embodied AI [9], [10]. My research on humane automation wishes to contribute to these goals by importantly and additionally implying a layer of consciousness regarding the social and labor context in which robotics and AI are to be situated. In other words, developing humane automation requires an in-depth, contextualized understanding of the context, the people, and their activities. Furthermore, the technical development of robotics and AI should not be shy to embrace the social and ethical specificity emerging from the context. As in my research, the proposed framework would need to be immersed in, inspired by, and eventually respond to the construction workers and their current activities. Therefore, documenting how technical decisions reflect these social grounds, and how social considerations affect automation and AI in meaningful ways, will be an inherent part of the proposed Ph.D. research.

Problem Setup

The intersection of three domains

The research is uniquely placed at the intersection of three domains: Artificial Intelligence (AI), robotics, and on-site construction. As a critical context, on-site construction is known for its labor intensity and lack of automation. With highly bespoke needs varying on a building-to-building basis, construction sites are fundamentally different from stipulated mass-production pipelines in organized factories. Existing approaches to on-site construction automation, which often reformulate a singular manual procedure with a task-specific autonomous system, still see limited success in dealing with real-world challenges such as on-site adaptability and system maintenance. Examples of these systems, such as brick-laying [11], spray-coating [12]–[14], ground-leveling [15], timber assembly [3], [16]–[18], rebar tying and handling [19], [20], seek to use robots to partially, or completely, replace existing human workers. However, few of them take advantage of, or build on, the tacit knowledge and composite skills performed by experienced human workers and assume an overly utopian experiment setting rather than an actual, messy construction site. Reflecting on these constraints, and drawing meaningful insights from the concrete site study conducted by the ReAC group, I wonder if there would be a more human-centered way of developing mobile robots that can blend into and support existing workflows, such that the latest advancement in AI and robotics can be effectively utilized in a grounded and practical manner. In other words, rather than simplifying and idealizing the real-world challenges, can the technical development of robots be meaningfully contextualized to the true "messiness" of construction contexts and activities? Can robotic advancements be leveraged toward a supportive, adaptive, and socially comfortable work capability?

Another defining piece of the puzzle is in the domain of artificial intelligence (AI). Reinforcement Learning (RL) is often considered a subfield of AI and an area of Machine Learning (ML) concerning how agents learn to act in environments by maximizing an expected cumulative reward [30]. Compared to other canonical methods, such as expert-supervised solutions or hand-crafted models, RL methods employ agents with a decision-making policy by learning through trial-and-error interactions in the environment. With the success of Deep Reinforcement Learning (DRL) in AlphaGo [32], there has been increasing interest in applying RL to various robotic control tasks such as contact-rich manipulation [33]–[37], humanoid bipedal walking [38]–[40], autonomous car driving [41], [42], and drone flight control [43], [44]. RL methods are useful in these tasks because real-world robotic behaviors often occur in dynamic, complex environments and are sophisticated for explicit modeling or hand detailing [31]. In recent years, RL research and applications have also emerged in the domain of construction. For instance, Apolinarska et al. [16] used a distributed RL workflow to automate timber joint assembly. Liang et al. [21] taught robots to perform quasi-repetitive construction tasks by observing human demonstrations. In previous work by the ReAC group [22], we proposed a DRL-driven multi-drone collision avoidance framework for the semi-autonomous bricklaying of bespoke structures. However, as a common trait, robots in these examples mostly interacted with materials or other robots. The coexistence of human workers in the same working environment was largely ignored, and the adaptability of RL was mostly employed to enhance robot-material or robot-robot dexterity. The potential of using RL for on-site mobility around workers, let alone providing meaningful support, is still largely an unexplored area. Considering the described hidden gap, a question to be asked is whether it would be possible to leverage the advantage of RL to enable adaptive robot navigation in the same dynamic indoor environment as busy construction workers, such that handy robot work assistance can be offered around them, even within fairly close range.

In the robotics domain, there is also a burgeoning trend to develop mobile robots that can operate in human-populated environments. The task of social robot navigation has undergone decades of research and development. First demonstrated in museum-guiding robots such as RHINO [23] and MINERVA [24], the attention of the area has gradually shifted from treating people as non-reactive, simple dynamic obstacles [25], [26] to lively, behavioral moving agents [27], [28]. The capabilities of robots have also expanded from plain collision avoidance to a richer body of social awareness and social compliance. For example, robots are expected to avoid interfering or cutting across existing social groups [29], approach people comfortably and politely [30], engage with people in human-like patterns by respecting interpersonal spaces [31], [32], and navigate efficiently while observing social dynamics [27], [33]. In addition to theoretical and algorithmic developments, practical robotic applications have also appeared in various domains, such as Ambient Assistive Living (AAL) and elderly caregiving [34], delivery and service in public spaces [35], [36], and interactive therapy and education for children [37]. In recent years, following the work of CADRL [20] and GA3C-CADRL [45], developed by a team of MIT researchers, a rich body of RL research has expanded the area in various aspects [19], [46], [47]. Compared to canonical methods, RL is particularly useful in training robots to adaptively predict and react to dynamic human crowds in an implicit, coupled manner [27]. Despite the promise, many RL-based methods assumed an open, spacious context such as outdoor campus space, open public squares, empty classrooms/labs, and wide corridors. With a few exceptions [38], limited attention has been paid to applying and examining these RL methods on messy, cluttered, and populated indoor construction floors. Witnessing the capabilities, as well as limitations, shown in these advancements, I wonder if it would be possible to bring the latest robotic flexibility, mobility, and interactivity to labor-intensive construction work contexts, where RL-based social navigation methods can potentially enable useful robotic engagement in workers' current practices.

Research question

Rather than entirely replacing humans, how can advancements in RL and robotics be leveraged to create robots that can adaptively support, accompany, and assist workers in their existing labor-intensive workflows, specifically in the on-site construction context?

Hypothesis

RL-based social navigation methods can work in concert with (1) robot hardware prototyping, (2) unstructured site perception and robot state estimation, (3) worker detection and tracking, (4) efficient and hierarchical motion planning, and (5) contextualized fine-tuning to collaboratively support on-site construction workers and realize navigation comfort and adaptability in effective, measurable ways. Synthesized as an RL-driven robotic framework, it should (1) be able to enable safe and fluent robot navigation on an existing complex construction site, populated with busy workers, (2) observe and adapt to workers' movements and activities such that close-range worker-robot interactions can proceed with certain degrees of social comfort, and (3) be able to offer meaningful forms of live work assistance that are helpful in a labor-intensive construction workflow.

Methods

In order to explore the above hypothesis, my research follows a 'research by prototyping' [39]–[41] approach. In other words, a contextualized robot prototype is to be designed, developed, and tested as a realization and validation of the proposed RL-driven robotic framework. Drawing from the ReAC group's concurrent qualitative study of a construction site, carpentry formwork is selected as a tentative robot support scenario for investigation. In this context, the prototype would roughly take the shape of a "work companion rover", a UGV (Unmanned Ground Vehicle) that can navigate autonomously and adaptively among customizable key locations alongside workers upon simple, intuitive commands. The robot intends to relieve workers from burdensome site traversals and weight-bearing by delivering and carrying essential yet heavy tools and materials nearby. The framework, as well as the proposed support functions, will be evaluated quantitatively and qualitatively with the robot prototype through both in-lab experiments and on-site demonstrations.

The proposed framework employs RL as the algorithmic backbone and consists of four key components: (1) robot system design and hardware setup, (2) RL infrastructure of site perception, worker detection, and robot state estimation (3) hierarchical and efficient motion planning, and (4) contextualized training and deployment of the RL algorithm. The framework is developed in tight conjunction with the technical research conducted in the ReAC[1] group. Through the two-year effort, I participated in an entire research cycle including proposal drafting, lab space and sandbox setup, technical roadmap drafting, robot prototype development, in-lab evaluations, and on-site demonstrations. I should also acknowledge the valuable contributions of fellow researchers[2] in helping with different aspects of the framework's development, and in contributing qualitative observations and analyses that crucially informed it.

Key elements

Below is a more detailed account of the four components of the proposed framework. Robot system design and hardware setup include the entire body of work related to robot system design, both on the hardware side and the software side. On the software side, the effort entails a systematic roadmap and modular breakdown (Fig. 2) for realizing the robot prototype. The hardware effort responds to software

---

[1] The Rethinking Automation in Construction (ReAC) group is an interdisciplinary team of researchers situated in the Computational Design Laboratory of Carnegie Mellon University. The group investigates how robotic systems can support humans in construction by engaging with actual construction sites and people. The current effort is formed by two highly intertwined components. The first component, qualitative research, is to understand existing construction labor, scenarios, and workflows through on-site ethnographic studies. The second component is technical research, which involves developing and testing an RL-driven robot prototype based on key insights drawn from the site study.

[2] On the technical side, Jiaying Wei offered critical help in the building, setup, and diagnosis of the robot's initial navigation stack, structures and sensors. She also participated in later works for PDT and RL integration. Yuchen (Joshua) Cao contributed to the component of LiDAR Odometry and Mapping. Sam Shun from NREC provided critical support in designing and implementing the initial LiDAR and IMU synchronization box for LOAM. Meghdeep Jana helped with RL algorithm selection, testing, and benchmarking during the project's early stages. Tingsong Ou also spared much time in robot tuning when the project was facing hardware issues in the early days. On the qualitative study side, the ethnographic research of the construction site and labor was conducted by Emek Erdolu, another PhD student in the ReAC group.

needs and consists of work related to (1) robot chassis and control, (2) robot structure, (3) onboard computing, (4) holistic sensor package, and (5) additional modifications. The building process starts from a blank 4WD Clearpath Husky robot base[3] and takes a rather incremental, "baby step" process. On countless occasions, new corner cases and issues may emerge from either site observation or lab testing. Modifications, upgrades, repairs, and customizations will become a weekly routine. To embrace the challenges, a general guideline behind hardware effort is rapid prototyping and adaptation. The final appearance of the robot is also an outcome of iterations of optimizations. In terms of individual hardware components, various considerations for on-site deployment need to be taken into account. For onboard computing, the major concern is the GPU acceleration for multiple real-time Deep Neural Network (DNN) inference. As for the sensor package, the focus is to gather data streams that can sufficiently and efficiently capture important features of the messy, complex surroundings. The robot structure and additional installations are closely related to performance robustness, comfortable human-robot interaction, and worker habits. Overall, the system design and hardware setup effort will lay a concrete foundation for the development of the framework.

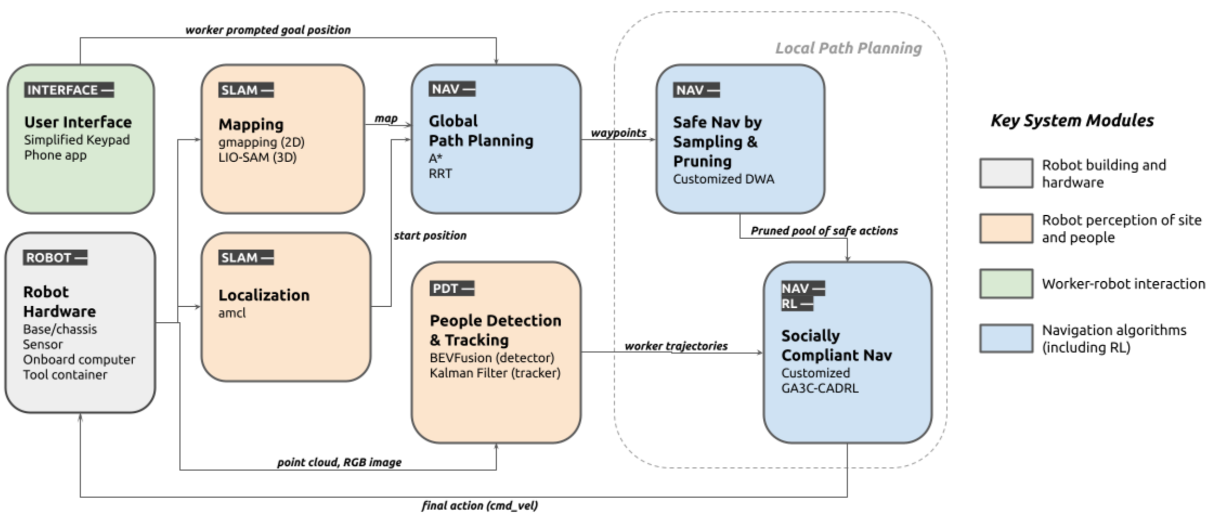

Figure 2. System design of the RL-driven robotic framework for RSCC. This figure is/will also be used in the CSCW paper publication (currently under review) [X]

The RL infrastructure effort is an essential prerequisite for the real-time execution of RL algorithms on a real robot. In virtual simulators, the robot's observable states can easily be extracted and supplied through simple numeric calculations. The case is no longer true in the real world. Every source of input needs to come from raw sensory data and will require additional processing. Generally speaking, the real-time input for RL includes (1) the position and velocity of the robot, (2) the position and velocity of every nearby person, (3) the relative position of the goal point, (4) the robot's current heading angle, and (5) the preferred speed of the robot. To be able to provide these data with low latency, three major components are at stake, namely (1) site perception and SLAM, (2) people detection and tracking (PDT), and (3) robot

[3] The robot has been kindly made available through a joint effort by Prof. Jean Oh, colleagues from the National Robotics and Engineering Center (NREC), and Sam Shun.

state estimation. Site perception and SLAM combines 2D (gmapping) and 3D (Octomap, LOAM) methods/tools to build an initial map of the site. It answers the question of what the construction site should look like for successful navigation and will be used as a prior for robot localization and relocalization. People detection and tracking uses both vision-based and LiDAR-based methods to offer real-time estimates of nearby workers' moving trajectories. Finally, the robot state estimation component employs LiDAR Odometry and Mapping (LOAM) [42] typed methods to produce a high-quality odometry source for robot localization, especially when solely relying on scanning-matching-based Monte-Carlo Localization (e.g. AMCL) methods would fail on non-stationary and unstructured construction floors. The entire infrastructure is organized and built upon ROS. Each component is made relatively modular such that the comparison, switching, and upgrading of individual methods become available efficiently and practically.

The hierarchical stack for motion planning details how RL is positioned in an efficient, composite on-site navigation strategy. As aforementioned in Section "Problem Setup", many promising RL methods in social robot navigation assume an open space context, thus they mainly consider people's trajectories when developing the algorithm, not essentially accounting for static obstacles (e.g. boxe, column, wall) in the experiment layout [43], [44]. In other words, the input of these RL algorithms only contains information about neighboring people, not static obstacles. On the other side, there has been an abundance of efficient methods specifically dealing with static obstacle avoidance. Drawing inspiration from previous research in the field [28], [45], I will approach the motion planning problem hierarchically. Instead of overloading the RL model with additional static obstacle avoidance capabilities, the composite motion planning is achieved in two steps. In the first step, I will mainly choose popular lightweight searching-based and sampling-based methods such as DWA, and RRT* for safe navigation. At this step, the robot is expected to quickly rule out motion primitives that could lead to any collision. Then in the second step, RL will decide among the remaining motion primitives the most socially compliant one that can adapt to nearby workers' moving trajectories. An advantage of this hierarchical approach is that it is efficient and generalizable among a pool of methods. Therefore, it will be feasible to compare the real-world performance of different classical and RL methods with minimal engineering efforts. To decrease latency in deployment, several measures will be necessary to accelerate RL inference and ensure RL and classical methods work in tandem seamlessly.

Finally, the contextualized training and deployment of RL aims to make the algorithm more situated and appropriate for the carpentry work scenario. Unlike scenarios commonly appearing in social navigation research, where the robot can operate among dense pedestrian crowds (10-15 people), construction sites are less populated (mostly less than 4 people). However, this feature does not necessarily simplify the problem. On the contrary, challenges arise in other aspects. The space for navigation is comparatively much more constrained, and the safety expectation when in close encounters with workers is higher. Carpentry workers, unlike pedestrians, are mostly packed with intense work at hand in tight spaces against the wall and do not have extra attention to spare for the robot. To accommodate the uniqueness, the algorithm trained using generic experimental setups in the simulator needs additional alignment and improvement before it can be applied to the robot. A curriculum of experiments will need to be manually designed in the symbolic simulator based on key site observations (Fig. 3). The maximum number of agents is set to be 4-6, and the free space around agents upon initialization is set to be narrower. Some close encounter cases need to be intentionally crafted to increase the difficulty. A pretrained RL model (e.g. CADRL, GA3C-CADRL, SARL) will need to be finetuned in courses using these additional setups. Eventually, the finetuned model, rather than the original vanilla model, will be used on the robot as the final inference engine.

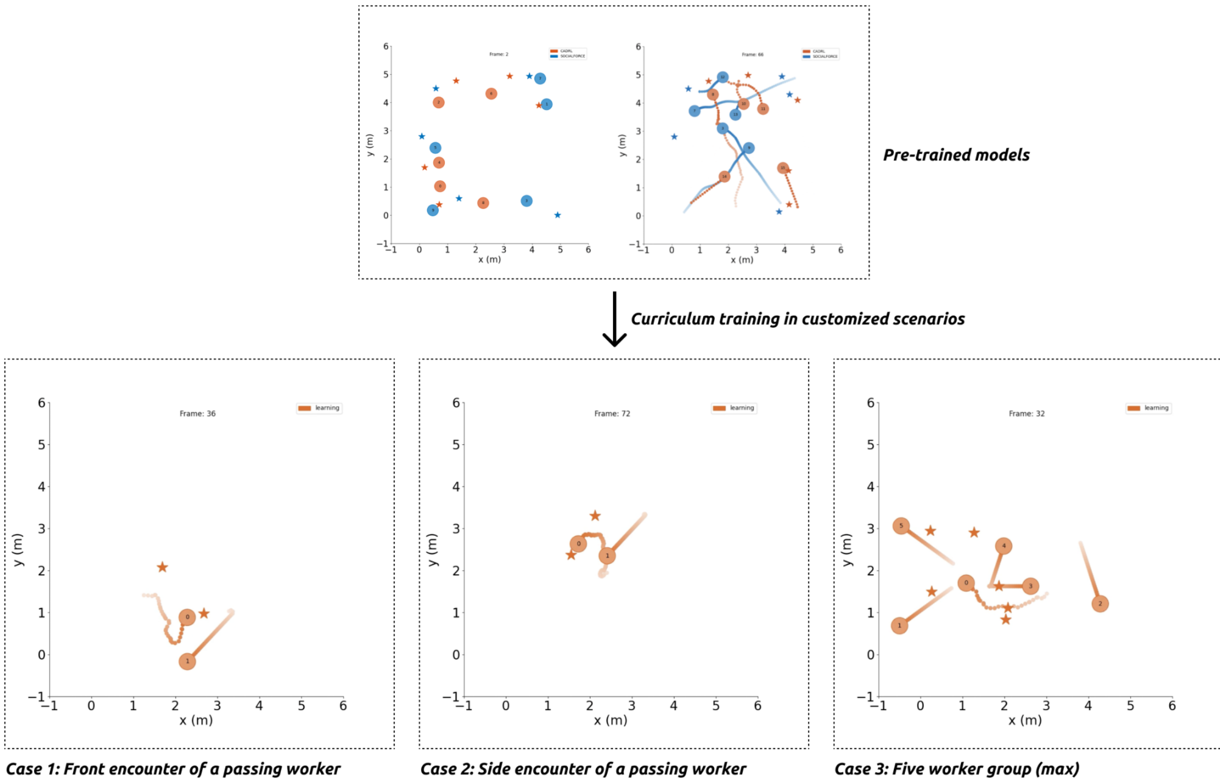

Figure 3. Additional curriculum training for contextualizing the RL algorithm for construction sites. This figure is/will also be used in the CSCW paper publication (currently under review) [X]

Evaluation and demonstration

The framework will be demonstrated and evaluated in two phases, namely the lab phase and the on-site phase. In the lab phase, the robot is to be tested extensively to (1) uncover potential technical issues, (2) test limits and extremes of navigation capabilities (e.g. maximum distance of traversal, corner cases, battery life, module latencies), and (3) quantitatively evaluate the framework's performance through an ablation study. More specifically, the robot's performance will be compared when the RL layer is enabled and disabled. A series of curated test cases will be designed and conducted in a sandbox environment in the lab, echoing representative scenarios observed on-site. The sandbox environment will be created using common construction objects and will aim to reproduce the layout features on real sites. In addition to common social navigation metrics, such as the rate of successful navigations/deliveries, collisions, and time-outs, two comfort-related metrics will be specifically adapted in our context to examine the framework's capability to offer comfortable support near workers. The first metric is the rate of comfortable encounters. Besides not bumping into workers, a comfortable encounter in our context entails that (1) the robot shall not be within 15cm of the worker while navigating (especially the knee and the foot), (2) the encountered worker should not be surprised by the robot, or need to split any attention to adjust to the robot, (3) the robot shall not interrupt any work process, and (4) the encountered worker generally considers the interaction to be comfortable. The second metric is the rate of comfortable delivery. Compared to successful delivery, a comfortable delivery is one that (1) does not take too long to finish/adjust, (2) ends at a position and orientation easily accessible to workers reaching out for tools/materials, (3) does not cut through any work process the worker is engaging in, and (4) does not lead to any complaint or concern by the worker.

In the on-site evaluation phase, more emphasis is placed on qualitatively demonstrating and assessing the framework with direct participation from actual carpentry workers on a real-world construction site. Particularly, three demo scripts are designed. In Script 1, the robot navigates between a work zone and a loading station upon workers' controller command, depicting the scenario of the robot assisting workers in picking up tools and hardware from a supply station. In Script 2, the robot migrates from one work zone to another alongside the workers while carrying heavy loads, demonstrating the tool/material carrying capability of work companions in procedural workflows. In Script 3, the robot navigates to a distant, unfamiliar random spot on site from the loading station upon command, stress-testing its navigation limit on complex and dynamic construction floors. In all three scripts, team members will intentionally walk on-site to test the robot's performance in social navigation and collision avoidance. Besides the three scripts, an interview will be conducted with participating workers to hear their candid thoughts about the proposed framework and potential future improvements.

Expected Contribution and Outcome

This research advances the state of the art in computational design by bringing together AI, robotics, and construction in a new way. In particular, it shows how the latest advancements in RL and social robot navigation can be mobilized toward a human-centered technical framework that can support on-site carpentry workers. In more technical detail, the research contributes,

1) A robot prototype demonstrating on actual construction sites how reinforcement learning algorithms can seamlessly and hierarchically work in conjunction with a group of supplementary methods to offer comfortable robot encounters and deliveries, which is central to supporting the daily work of carpentry workers.
2) A generalizable, efficient RL-driven robotic framework, including both hardware and software effort, that can serve as a navigation base for enabling flexible and mobile robot support functions in indoor, cluttered, labor-intensive work scenarios.
    a) An effective pipeline for employing, contextualizing, and finetuning a group of recent RL-based social navigation algorithms for empowering human-robot interaction in a specific work context.
    b) A modular and open-source design and implementation of robot system and hardware with a complete stack of robot control and sensory packages geared towards rusty indoor work environments.
    c) A robust architecture of perception modules tailored to physically complex, unstructured obstacles and the efficient detection of workers' activities.
    d) A lightweight and hierarchical navigation stack for seamlessly integrating RL algorithms with canonical collision avoidance methods.
3) A case study of a technology developed in close dialogue with the socio-technical context of on-site construction workers and their actual daily workflows.
4) A concrete robotic foundation for future investigation in algorithmic, ethical, social, and applicational aspects of humane automation and AI.

At a higher level, I wish that the proposed framework can advocate and substantiate my vision regarding "Humane Automation". I certainly wish the image of robots collaboratively and sensitively supporting and working with people, nicely depicted with TARS in "Interstellar", MB-26 Benben in "The Wondering Earth II", and Baymax in "Big Hero 6", could be realized in the near future. Coming from an architecture background and witnessing the long-existing difficulties occurring in highly manual industries, I choose carpentry work as a demonstrative entry point. The research offers a concrete case study where robots humanely help and accompany people in existing laborious daily work. The research also provides a solid and reusable technical foundation in wish to trigger more contextualized and considerate research in labor-intensive work scenarios, towards a robotically supported collaborative future.

In addition to these contributions, I also wish to candidly share the tacit knowledge, difficulties, failures, and successes encountered during the research process, in the hope of triggering a deeper reflection of AI and robotics research and development in general. As the current norm, most published papers often only document the final technical choices, that is, what equipment is used, what algorithm is selected, and what procedures are performed. Although I enjoy the clarity and brevity, they largely miss the actual process of solving real-world problems, which inevitably involves iterations of twists and turns, trials and errors.

Many of the seemingly trivial issues can have an unexpected, long-tailed, even defining effect on the entire system. I mostly manage to resolve these issues by watching numerous YouTube tutorials, reading countless blog posts, and scrutinizing pages of GitHub Issues. At that time, I often hoped there could be a more explicit way of sharing those practical and critical experiences. Therefore, I wish an honest, reflective, systematic documentation of each phase of the research can benefit researchers facing similar difficulties, and for the community at large.

Timeline

2019.8 - 2021.8	Coursework and preliminary research
2021.8 - 2022.8	Qualitative site research Technical development of robot prototype - phase I - Basic hardware structure and design (sensor, computer, chassis) - Baseline sampling-based navigation stack - Site perception and LOAM - RL pre-training and curriculum training in the simulator
2022.8 - 2022.9	Lab demonstration - phase I - Sampling-based navigation stack - Basic collision avoidance among walking people - Long-distance navigation in spacious environments
2022.9 - 2023.4	Technical development of robot prototype - phase II - Efficient people detection and tracking stack - Hardware upgrades and tuning - RL integration with upstream threads - System optimization and improving computational efficiency
2023.4 - 2023.5	Intensive lab test and debugging Lab demonstration - phase II - RL-driven social navigation stack in complex environments - People detection and tracking - Human-robot interaction and long-distance summoning - Battery limit test Preparation for the on-site demo
2023.5	On-site demonstration - Script 1, navigate between the main station and a work zone - Script 2, migrate from workzone 1 to workzone 2 - Script 3, navigate to an unfamiliar, distant position Interview with workers
2023.5 - 2023.9	Qualitative evaluation and analysis of gathered data, In-lab experiment for quantitative evaluation of robot support Paper submission to ACM SIGCHI CSCW after major revision
2023.11 - 2024.8[4]	Focused thesis writing
2023.12	Proposal
2024.3	Complete thesis draft
2024.8	Thesis submission and defense

[4] This is the current progress (where I am with the research)

Previous research

As critical preliminaries of the current proposal, I should also mention some of my previous work, and how they have critically inspired and paved the way for the current research. In a previous work of the ReAC group, the team and I worked on an RL-driven, multi-agent collision avoidance framework that can enable brick-laying drones to build flexible structures [22]. The research raises a novel robot-centric autonomous system, where material pickup, charging, and site setup are all subtly redesigned in the simulation. The role of humans is mostly placed in maintenance, supply, and handling which is essential but difficult to automate. By design, the drones are not to work alongside humans, and any manual intervention would require complete pausing of the system. The research on the one hand shows how capable RL algorithms can be in dealing with complex, dynamic collision avoidance scenarios; on the other hand, also shows how incapable the system can be in handling real-world messiness, trivialities, and complexities. The project essentially assumes a "human-in-the-loop" approach. Like its many predecessors, the project inserts the latest AI methods into a specific, idealized link of an otherwise complex system, and expects humans to compensate for the unspoken, undocumented, insignificant rest. Reflecting on the limitations of this approach led me to where I am now. Inherently, the current proposal exemplifies a contrary "robot-in-the-loop" approach where the significance of humans in bespoke, labor-intensive trades is no longer intentionally downplayed.

In another research with Autodesk AI Lab and Robotics Lab, I studied how rewards can affect the training and performance of RL in complex real-world robotic tasks. Unlike playing a video game where rewards are easily pre-defined, real-world tasks such as robotic manipulation and social robot navigation often require extensive expert knowledge in designing the reward function. My research shows that instead of hand-crafting rewards based on experience, it is possible to effectively train the robot using a learned reward function through self-supervised learning. Though not directly applied in the current research, the study benefits the current research by showing how tuning weights of the reward component may affect the learning results.

Annotated Table of Chapters

1. **Introduction**
   This chapter will give an overview of the research, including its context, method, experimentation, and potential social and technical impact.

2. **Background**
   a. **On-site construction automation**
   b. **Reinforcement Learning (RL) methods**
   c. **Social robot navigation**
   (Drafting in progress)
   The background chapter will give a literature review of the three disciplinary components of the research, namely AI/RL, robotics, and on-site construction automation. In each section, the current advances in each field related to the proposed research, as well as potential gaps and opportunities, will be discussed.

3. **Context of on-site carpentry formwork**
   (Drafting in progress)
   This chapter will specifically zoom in on the context of on-site carpentry formwork as a scenario for study. It is an extension to the ReAC group qualitative study effort. In a way that rather than providing a full ethnographic account of the site observation, it will focus on the robot support scenarios and key robotic functions distilled from the observations. It also covers technical attributes and difficulties arises from these scenarios that will serve as a critical prior for the development of the RL-driven robotic framework.

4. **Hypothesis**
   (Drafting in progress)
   This chapter will raise the question and hypothesis of the research.

5. **Methods**
   (Completed 1st draft)
   This chapter will describe in detail the methodology used to enable the proposed framework. It will include both hardware and software efforts. An account of the system design and its key components will be provided, including the reasons and iterations.
   a. **Robot system design and hardware setup**
   This section will describe the hardware formulation and structure of the robot.
   b. **Robotic perception of the site and people**
   The real-time inference of RL requires several processed inputs from raw sensory data streams. These include robot state estimation, the relative position of the goal, and neighboring workers' trajectories. This section will describe the selected solutions for the software infrastructure of RL, and will also give a brief account of the encountered iterations and trial-and-errors in forming these solutions.
   c. **Efficient and hierarchical motion planning**

The employed RL method is mostly used for social robot navigation in open space. However, the construction site is cluttered with all kinds of unstructured obstacles. Instead of overloading the RL algorithm with the capability of both static collision avoidance and social navigation, the framework implemented a hierarchical motion planning stack such that a canonical searching-based method can work in tandem with the RL algorithm.

d. **Contextualize RL-driven social navigation algorithm in carpentry work scenarios**
This section details how RL-driven social navigation algorithms can be employed, contextualized, and finetuned in response to the specific work context. Besides generic training in the simulation environment, a series of designed experiments echoing workers' activities are used to curriculum-train the algorithm, such that the commonly observed close worker-robot encounter cases can be addressed in the training stage before deploying to the physical robot.

6. **Results**
(Drafting in progress)
This chapter will offer a detailed description of the research results from different perspectives.

a. **Robot prototype**
This section will describe in detail the final outcome of the robot prototype, including the chassis setup, sensors, structures, and worker-oriented design.

b. **On-site demonstration**
This section will document the on-site demonstration and evaluation. The robot is tested on a real-world construction site with workers involved in three highly interactive support cases.

c. **Lab-based quantitative evaluation**
To quantitatively measure the effectiveness of employing RL for robot support and work companions, a sandbox environment is set up in the lab echoing layouts on-site. An ablation study will be documented to numerically compare the robot's performance when enabling and disabling the RL layer.

d. **Worker interviews**
This section will document the feedback and insights drawn from the short interview conducted among participating workers.

7. **Conclusions**
This chapter will summarize the research with key findings and conclusions.

8. **Reflections**
This chapter will serve as a reflection of the research, critically discussing its resolved and unresolved issues.

9. **Contributions**
This chapter will concisely summarize the contributions of this research, especially in the domain of computational design, robotics, and human-centered AI.

10. **Future Work and Next Steps**

This chapter will briefly outline the potential RL algorithmic work as a follow-up of the current robotic framework. It will be a critical next step towards the goal of humane automation and AI.

Bibliography

[1] A. Menges, “Fusing the Computational and the Physical: Towards a Novel Material Culture,” *Archit. Des.*, vol. 85, no. 5, pp. 8–15, 2015, doi: 10.1002/ad.1947.

[2] G. Retsin, “Bits and Pieces: Digital Assemblies: From Craft to Automation,” *Archit. Des.*, vol. 89, no. 2, pp. 38–45, 2019, doi: 10.1002/ad.2410.

[3] V. Helm, M. Knauss, T. Kohlhammer, F. Gramazio, and M. Kohler, “Additive robotic fabrication of complex timber structures,” 2016, pp. 29–44. doi: 10.4324/9781315678825-3.

[4] J. Burger, T. Huber, E. Lloret-Fritschi, J. Mata-Falcón, F. Gramazio, and M. Kohler, “Design and fabrication of optimised ribbed concrete floor slabs using large scale 3D printed formwork,” *Autom. Constr.*, vol. 144, p. 104599, 2022.

[5] B. Shneiderman, *Human-centered AI*. Oxford University Press, 2022.

[6] W. Xu, “Toward human-centered AI: a perspective from human-computer interaction,” *interactions*, vol. 26, no. 4, pp. 42–46, 2019.

[7] T. B. Sheridan, “Human centered automation: oxymoron or common sense?,” in *1995 IEEE international conference on systems, man and cybernetics. Intelligent systems for the 21st century*, IEEE, 1995, pp. 823–828.

[8] J. E. Michaelis, A. Siebert-Evenstone, D. W. Shaffer, and B. Mutlu, “Collaborative or Simply Uncaged? Understanding Human-Cobot Interactions in Automation,” in *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, in CHI ’20. New York, NY, USA: Association for Computing Machinery, Apr. 2020, pp. 1–12. doi: 10.1145/3313831.3376547.

[9] J. Duan, S. Yu, H. L. Tan, H. Zhu, and C. Tan, “A survey of embodied ai: From simulators to research tasks,” *IEEE Trans. Emerg. Top. Comput. Intell.*, vol. 6, no. 2, pp. 230–244, 2022.

[10] M. Savva *et al.*, “Habitat: A platform for embodied ai research,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 9339–9347.

[11] T. Bonwetsch, “Robotically assembled brickwork: Manipulating assembly processes of discrete elements,” Doctoral Thesis, ETH Zurich, 2015. doi: 10.3929/ethz-a-010602028.

[12] J. McAlorum, M. Perry, C. Vlachakis, L. Biondi, and B. Lavoie, “Robotic spray coating of self-sensing metakaolin geopolymer for concrete monitoring,” *Autom. Constr.*, vol. 121, p. 103415, Jan. 2021, doi: 10.1016/j.autcon.2020.103415.

[13] S. Chaltiel, M. Bravo, S. Goessens, P. Latteur, M. Mansouri, and I. Ahmad, “Dry and Liquid clay mix drone spraying for Bioshotcrete,” presented at the Proceedings of IASS Annual Symposia, International Association for Shell and Spatial Structures (IASS), 2018, pp. 1–8.

[14] A. Uryasheva, M. Kulbeda, N. Rodichenko, and D. Tsetserukou, “DroneGraffiti: autonomous multi-UAV spray painting,” in *ACM SIGGRAPH 2019 Studio*, in SIGGRAPH ’19. New York, NY, USA: Association for Computing Machinery, Jul. 2019, pp. 1–2. doi: 10.1145/3306306.3328000.

[15] L. Yang, Z. Huixing, and W. Shun, “Design and Simulation of Ground Leveling Construction Robot.,” *J. Beijing Inst. Civ. Eng. Archit.*, vol. 37, no. 2, 2021.

[16] A. A. Apolinarska *et al.*, “Robotic assembly of timber joints using reinforcement learning,” *Autom. Constr.*, vol. 125, p. 103569, May 2021, doi: 10.1016/j.autcon.2021.103569.

[17] S. Leder, R. Weber, D. Wood, O. Bucklin, and A. Menges, “Distributed Robotic Timber Construction,” p. 10.

[18] A. Kunic, R. Naboni, A. Kramberger, and C. Schlette, “Design and assembly automation of the Robotic Reversible Timber Beam,” *Autom. Constr.*, vol. 123, p. 103531, Mar. 2021, doi: 10.1016/j.autcon.2020.103531.

[19] C. A. Cardno, “Robotic Rebar-Tying System Uses Artificial Intelligence,” *Civ. Eng. Mag. Arch.*, vol. 88, no. 1, pp. 38–39, Jan. 2018, doi: 10.1061/ciegag.0001260.

[20] J. Jin *et al.*, “Robotic binding of rebar based on active perception and planning,” *Autom. Constr.*, vol. 132, p. 103939, Dec. 2021, doi: 10.1016/j.autcon.2021.103939.

[21] C.-J. Liang, V. R. Kamat, and C. C. Menassa, “Teaching robots to perform quasi-repetitive

construction tasks through human demonstration," *Autom. Constr.*, vol. 120, p. 103370, Dec. 2020, doi: 10.1016/j.autcon.2020.103370.

[22] 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 *the 40th Annual Conference of the Association for Computer-Aided Design in Architecture*, 2020.

[23] W. Burgard *et al.*, "Experiences with an interactive museum tour-guide robot," *Artif. Intell.*, vol. 114, no. 1, pp. 3–55, Oct. 1999, doi: 10.1016/S0004-3702(99)00070-3.

[24] S. Thrun *et al.*, "MINERVA: A second-generation museum tour-guide robot," in *Proceedings 1999 IEEE international conference on robotics and automation (cat. No. 99CH36288C)*, 1999.

[25] D. Fox, W. Burgard, S. Thrun, and A. B. Cremers, "A hybrid collision avoidance method for mobile robots," in *Proceedings. 1998 IEEE international conference on robotics and automation (cat. No. 98CH36146)*, IEEE, 1998, pp. 1238–1243.

[26] M. Montemerlo, N. Roy, and S. Thrun, "Perspectives on standardization in mobile robot programming: The Carnegie Mellon navigation (CARMEN) toolkit," in *Proceedings 2003 IEEE/RSJ international conference on intelligent robots and systems (IROS 2003)(Cat. No. 03CH37453)*, IEEE, 2003, pp. 2436–2441.

[27] C. Mavrogiannis *et al.*, "Core Challenges of Social Robot Navigation: A Survey." arXiv, Mar. 16, 2021. doi: 10.48550/arXiv.2103.05668.

[28] 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, Sep. 2018. [Online]. Available: https://arxiv.org/pdf/1805.01956.pdf

[29] K. Katyal *et al.*, "Learning a group-aware policy for robot navigation," in *2022 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, 2022, pp. 11328–11335. doi: 10.1109/IROS47612.2022.9981183.

[30] Y. Kato, T. Kanda, and H. Ishiguro, "May I help you? Design of human-like polite approaching behavior," in *Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction*, in HRI '15. New York, NY, USA: Association for Computing Machinery, 2015, pp. 35–42. doi: 10.1145/2696454.2696463.

[31] P. Althaus, H. Ishiguro, T. Kanda, T. Miyashita, and H. I. Christensen, "Navigation for human-robot interaction tasks," in *IEEE international conference on robotics and automation, 2004. Proceedings. ICRA '04. 2004*, 2004, pp. 1894-1900 Vol.2. doi: 10.1109/ROBOT.2004.1308100.

[32] E. Torta, R. H. Cuijpers, and J. F. Juola, "Design of a Parametric Model of Personal Space for Robotic Social Navigation," *Int. J. Soc. Robot.*, vol. 5, no. 3, pp. 357–365, Aug. 2013, doi: 10.1007/s12369-013-0188-9.

[33] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation," *Int. J. Robot. Res.*, vol. 34, no. 3, pp. 335–356, Mar. 2015, doi: 10.1177/0278364914557874.

[34] L. V. Calderita, A. Vega, S. Barroso-Ramírez, P. Bustos, and P. Núñez, "Designing a cyber-physical system for ambient assisted living: A use-case analysis for social robot navigation in caregiving centers," *Sensors*, vol. 20, no. 14, 2020, doi: 10.3390/s20144005.

[35] J. Lee *et al.*, "ODS-Bot: Mobile robot navigation for outdoor delivery services," *IEEE Access Pract. Innov. Open Solut.*, vol. 10, pp. 107250–107258, 2022, doi: 10.1109/ACCESS.2022.3212768.

[36] M. Wen, J. Zhang, T. Chen, G. Peng, T. Chia, and Y. Ma, "Vision based sidewalk navigation for last-mile delivery robot," in *2022 17th international conference on control, automation, robotics and vision (ICARCV)*, IEEE, 2022, pp. 249–254.

[37] 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 Rehabil. Assist. Technol.*, vol. 1, no. 1, p. e1, Oct. 2014, doi: 10.2196/rehab.3151.

[38] Y. S. Kim, B. Lee, R. Murphy, and C. R. Ahn, “Context-appropriate Social Navigation in Various Density Construction Environment using Reinforcement Learning,” in *ISARC. Proceedings of the International Symposium on Automation and Robotics in Construction*, Waterloo, Canada: IAARC Publications, 2021, pp. 505–512. Accessed: Dec. 04, 2023. [Online]. Available: https://www.proquest.com/docview/2642454765/abstract/E836BA1EE767498EPQ/1
[39] R. Budde, K. Kautz, K. Kuhlenkamp, and H. Züllighoven, “What is prototyping?,” *Inf. Technol. People*, vol. 6, no. 2/3, pp. 89–95, 1990.
[40] X. Yan and P. Gu, “A review of rapid prototyping technologies and systems,” *Comput.-Aided Des.*, vol. 28, no. 4, pp. 307–318, 1996.
[41] I. Koskinen and J. Frens, “Research Prototypes,” *Arch. Des. Res.*, vol. 30, no. 3, pp. 5–14, Aug. 2017, doi: 10.15187/adr.2017.08.30.3.5.
[42] J. Zhang and S. Singh, “LOAM: Lidar Odometry and Mapping in Real-time,” in *Robotics: Science and Systems X*, Robotics: Science and Systems Foundation, Jul. 2014. doi: 10.15607/RSS.2014.X.007.
[43] 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)*, 2019, pp. 6015–6022.
[44] “Social Navigation Simulator (SNS).” cmubig, Feb. 25, 2022. Accessed: May 16, 2022. [Online]. Available: https://github.com/cmubig/Social-Navigation-Simulator
[45] M. F. Everett, “Robot designed for socially acceptable navigation,” Thesis, Massachusetts Institute of Technology, 2017. Accessed: Dec. 24, 2021. [Online]. Available: https://dspace.mit.edu/handle/1721.1/111698