Introducing Mobile Robots on the Automotive Final Assembly Line: Control, Sensing and Human-Robot Interaction

by

Vaibhav Vasant Unhelkar

B.Tech., Indian Institute of Technology Bombay (2012)
M.Tech., Indian Institute of Technology Bombay (2012)

Submitted to the Department of Aeronautics and Astronautics in partial fulfillment of the requirements for the degree of

Master of Science in Aeronautics and Astronautics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

Signature redacted

Author ....
Department of Aeronautics and Astronautics
May 21, 2015

Signature redacted

Certified by........... ....
Julie A. Shah
Assistant Professor
Thesis Supervisor

Signature redacted

Accepted by ....
Paulo Lozano
Chairman, Department Committee on Graduate Theses

Introducing Mobile Robots on the Automotive Final Assembly Line: Control, Sensing and Human-Robot Interaction

by

Vaibhav Vasant Unhelkar

Submitted to the Department of Aeronautics and Astronautics
on May 21, 2015, in partial fulfillment of the
requirements for the degree of
Master of Science in Aeronautics and Astronautics

Abstract

Traditionally, robots in manufacturing have been deployed in caged, static and predictable environments. Advances in robotics are enabling industrial robots to emerge from these traditional habitats, and enter the final assembly to work along side humans. My thesis contributes to this effort through development of a mobile robot capable of operating on final automotive assembly lines to assist humans.

Several algorithmic as well as design challenges exist when mobile robots enter the unpredictable, human-centric and time-critical environment of final assembly. My primary focus is on achieving autonomous mobility, a precursor for introducing robots to operational factory floors. Automotive assembly lines present a distinct challenge in form of surfaces that are dynamic, i.e., the conveyor belts which ferry cars in the factory. I develop a control strategy to enable autonomous navigation on such dynamic surfaces, and design a sensing module capable of detecting the conveyor belts. The designed system is tested in simulation, implemented on hardware and evaluated on an operational, automotive factory floor.

Evaluation in factory establishes preliminary success in the designed robotic system. Interesting, qualitative observations while introducing a robot in a real environment also emerge, and motivate need for enhancing the interaction capability of robots for time-critical tasks in human-centric environments. Towards this, we carry out a human subject experiment ($N = 24$) comparing the performance of the robot to that of a human assistant in an analogue assembly line environment. Results from the experiment provide a better understanding of the factors that impact fluency of interaction and inform the design of a more effective mobile robotic assistant. This work introduces mobile robots on the automotive assembly lines right next to people, thereby paving the way for utilizing them to assist busy, human associates in the myriad tasks involved in final assembly of cars.

Thesis Supervisor: Julie A. Shah
Title: Assistant Professor

Acknowledgments

First, I would like to thank Professor Julie Shah. Prof. Shah is an amazing mentor and truly an inspiration. Her technical advice and expertise have strongly shaped this research, and she has been a constant source of encouragement. Along with helping me grow as a researcher, she also ensured that I had a good work-life balance during the course of this project.

Designing a robotic system is a collaborative effort. In particular, I acknowledge my collaborators Johannes Bix, Dr. Alexander Bubeck and Dr. James C. Boerkoel Jr. who provided constructive feedback and lent their domain expertise. During this project, I also worked with two extraordinary undergraduate researchers, Jorge Perez and Ho Chit Siu, whom I would like to thank. I acknowledge Ron Wiken for his help with construction of the test environment.

The last two years would not have been as enjoyable without the continued support of my friends. I would like to thank my peers at the Interactive Robotics Group for the interesting discussions on topics ranging from robotics, pronunciation and life. I thank my friends at Tang, 88 Columbia, Edgerton, Sid-Pac, Cambridge East and beyond for the fun evenings involving board games, cooking and music. I also thank Divya for her support especially during the last stages of the thesis.

I am grateful to Beth Marois and Julia Finn of the Department of Aeronautics and Astronautics, and Dean Danielle Guichard-Ashbrook of the the International Students Office for supporting my stay at MIT. I would like to thank the comedians, and especially Mr. Letterman, for making all of us laugh. Lastly, I thank my parents, grandparents and sister for their unconditional care and love despite me being over a thousand miles away!

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Automotive industries have been one of the first to introduce robotics in their manufacturing processes. Indeed, it is no surprise that around fifty percent of the manufacturing in a typical factory today is done by industrial robots [1]. Figure 1-1a shows an example of one such completely automated system being used in the body shop of an automotive factory. However, driven by the requirement of high reliability, most of these industrial robots are caged, static and non-interactive. By operating in the cages and away from humans, the robot environment is rendered highly predictable allowing for both reliable and human-safe execution of largely pre-planned tasks.

Robots that operate beside or cooperatively with humans are envisioned as the "next generation of robotics" [8]. There exists an increasing demand to incorporate mobile interactive robots to assist humans in repetitive, non-value added tasks in the manufacturing domain. Due to the advancements in the state-of-the-art perception, planning, control and actuation, the boundaries for robots in manufacturing are being pushed by roboticists to bring them out of the cages and introduce them into final assembly along side humans. Krüger et al. [46] provide a detailed survey of robots being developed for and/or used in assembly lines. Several robotic aids [10] and collaborative robots [16, 57] have been developed as research prototypes as well as operational systems for final assembly, including automotive guided vehicles for delivering parts (Fig. 1-1b). However, till date there have been no mobile robots that can work along side humans on the automotive final assembly lines.

(a) Large, industrial robots performing the initial phase of car manufacturing. [5]

(b) Automated guided vehicles delivering parts in an automotive assembly line. [3]

Figure 1-1: Robots being used in different stages of car manufacturing in an automotive factory. (left) Large, industrial robots work in a caged setup away from humans to manufacture the car chassis through an 100% automated process. (right) Mobile robots, navigating using magnetic strips in the floor, deliver parts across large factory floors to human associates.

Our work is aimed at facilitating close-proximity, human-robot collaboration in automotive manufacturing by introducing mobile robots on the automotive final assembly lines right next to people. To work in close proximity with humans, there is a need for robots that can operate on the assembly lines and can navigate freely instead of being restricted to pre-decided lines on the factory floor.

This thesis presents a mobile robot capable of operating on automotive final assembly lines along side human associates, with focus on the design of motion control, sensing and human-robot interaction. The designed system is tested in simulation, implemented on hardware and demonstrated on an operational, automotive factory floor. The mobile robotic system can more directly assist humans working on the car during its final assembly, and allow for more flexible manufacturing processes. We believe research towards this robotic assistant, through the development of algorithms, realization of design guidelines and study of human-robot interaction experience, will not only introduce industrial robots on the automotive assembly floors but also aid in furthering robots that work with humans in other time-critical domains, including, factories, homes and roads.

In this chapter, we begin with a description of automotive final assembly lines including aspects that are relevant to design of a mobile robotic system. This motivates the key research challenges in introducing mobile robots on the automotive factory floor right next to people. Next, we define the scope of the thesis by describing the specific problems that we explored and the corresponding assumptions. Key contributions of the thesis are highlighted in Sec. 1.2.2. The chapter concludes with an overview of prior work on developing collaborative robots, with focus on robots developed for final assembly.

1.1 Automotive Final Assembly Lines

To accomplish our objective of developing a mobile robot for automotive factories, it is important to understand the characteristics of the assembly line environment and their implications on operation of a mobile robot. In this section, we describe the key features of the automotive assembly line environment, along with the key challenges it presents for mobile robots.

Figure 1-2 shows the typical, final assembly environment of an automotive factory. The environment is highly dynamic, uncertain and cluttered due to presence of human agents and mobile objects such as, cars on the assembly line and pick carts. In addition, some modern factories include line-following mobile robots used for delivery of parts across the factory floor [3]. Motion of human agents is unconstrained and they are free to move along the factory, both on and off the assembly line. Motion of other dynamic objects is also uncertain, and depends on the factory operations on a particular day, personal preferences, and unplanned events.

Along with the dynamic environment, automotive assembly lines have the distinct property of having surfaces that are dynamic, i.e., the conveyor belts which ferry cars in the factory. The motion of these conveyor belts and consequently that of the cars that they ferry is usually periodic; however, the conveyor belt can arbitrarily stop due to some event on the factory floor. While in motion the conveyor belts tend to move at a fairly slow (< 0.2 m/s) and almost constant speed.

Figure 1-2: A typical, final automotive assembly line [10]. The assembly line includes human associates, dynamic objects such as part-carts, and a conveyor belt which ferries the cars that are being assembled. The human associates primarily work on the conveyor belt, and periodically ferry tools and parts from outside the conveyor belt to carry out the car assembly.

The slow, yet periodic, motion of the conveyor belt results in each car spending typically less than three minutes at a workstation. Unscheduled stops in the assembly line are strict no-no, as they have high economic costs and human factor issues associated with them. This renders the tasks to be performed during final assembly highly time-critical, where every second matters. Any delays caused in car assembly may have cascading effects leading to disruption and/or inefficiencies in the assembly process. Any agent, including any prospective mobile robots, involved in car assembly should be able to successfully and repetitively complete its task within these short cycle times, as any incomplete work will lead to the undesirable, unscheduled stops. Presence of human agents further makes the task safety-critical, as safety is of paramount importance.

1.1.1 Implications for Mobile Robots

As described above, the automotive final assembly presents a highly dynamic and uncertain requirement that includes multiple human agents. This section describes the implications of such a challenging environment for design of a mobile robot that were derived based on observation of operations in operational automotive factory floors and discussion with domain experts.

Ensuring Human Safety

First and foremost, the robot should have mechanism both at the software and hardware level to ensure safety of humans and other equipment in the surrounding. This will require detecting obstacles/humans present in the surrounding, and modifying the robot behavior/plans accordingly.

Autonomous Navigation on Dynamic Surfaces

Next, prior to accomplishing any task on the automotive final assembly line, the robot first needs to be able to navigate autonomously on dynamic surfaces, i.e., the moving surfaces of assembly line. Achieving this capability requires the robot to fulfill multiple requirements that are described as follows,

- Sense obstacles, maintain a map of the environment, and estimate robot state: This should be done in presence of the dynamic surface, which can affect sensor performance, as well as dynamic objects and humans.

- Sense the location and state of the conveyor belt: The robot should have the capability to detect the state (speed) of the dynamic surface, as this information is not readily available to the robot in an automotive factory. Further, the robot should have the ability to detect the location of assembly line. Detection whether the robot is on the assembly line is critical, since even small errors/misdetection can be detrimental to robot's hardware and function.

- Trajectory tracking on dynamic surfaces: Mobile robots, using existing control algorithms, can track a desired plan generated by a user or path planners. However, to operate in automotive final assembly lines the mobile robot should be able to control its motion and track desired trajectories seamlessly, irrespective of whether it is on static surface, moving surface, or is straddling the assembly line (e.g., in the case when the robot is partially on the moving assembly line and remains partially on the static surface).

- Generating efficient path plans: The robot should be able to generate paths which are efficient, adhere to the time constraints of the task, and are capable to plan for goal locations that are moving (such as, a moving human that is to be assisted and/or car which is being assembled). This may additionally require tracking dynamic goals, such as, tracking the car/s being assembled.

Fluent Human-Robot Interaction

Along with being safe and navigating autonomously, the robot should also perform tasks efficiently and collaborate fluently with humans. To provide benefit, similar to other human-robot interaction scenarios, the robot should

- Perform deliberative task planning and scheduling: To function in dynamic manufacturing environments, the robot should be able to dynamically plan tasks, and maintain flexible schedule [20, 30].

- Infer the task/motion intent of human agents: Autonomous planning in dynamic environments can benefit if high fidelity, predictive information regarding the future state of the environment is available. By inferring the task/motion of human agents the robot can better reason for the dynamic, uncertain interactions in the final assembly.

- Convey its intent: Awareness of robot's intent will help the humans better reason and adapt while working with a robot, resulting in a more fluent collaboration.

Considerations for Factory Deployment

Issues concerning factory deployment of the robot should also be considered, as much as possible, in the design process of the robot. These requirements primarily include design of a system that

- requires minimal modifications to the factory layout and infrastructure,

- is easy to use and maintain, and

- includes redundancy for a more robust autonomous operation.

1.2 Scope of the Thesis

As discussed in the previous section, several algorithmic as well as design challenges exist when mobile robots enter the unpredictable, human-centric and time-critical environment of final assembly. This section describes

(i) the problem statements that we tackle in this thesis towards the multi-disciplinary challenge of introducing a mobile robot in automotive final assembly,

(ii) the robotic platform selected to demonstrate the designed capabilities, and

(iii) the outline of the thesis.

Autonomous Mobility on Conveyor Belts: Control and Sensing

In this thesis, firstly, we focus on *achieving autonomous mobility on conveyor belts*, a precursor for introducing robots to operational factory floors. Specifically, we undertake the design of a control and sensing module for robot navigation on dynamic surfaces. The goal is to demonstrate the designed autonomous mobility capability in an operational factory floor. Development of this enabling capability will allow for mobile robots to be used in automotive factory floors, thereby paving the way for utilizing them to assist busy, human associates in the myriad tasks involved in final assembly of cars that require autonomous mobility.

Human-Robot Interaction in Final Assembly

Irrespective of the specific task that the mobile robot carries out in the final assembly, it will always be working with or along side human associates. Small deficiencies in the human-robot interaction in the time-critical domain of automotive final assembly can significantly degrade the efficiency of overall work-flow. Hence, secondly, we seek to gain a better understanding of the factors that impact the *saliency and collaborative fluency* of a robot while working with a human during a time-critical assembly task.

1.2.1 Hardware Platform : Rob@Work 3

Design of a robotic system for automotive final assembly is heavily influenced by its hardware capabilities, for instance, differentially driven robots are not suitable as they cannot hold their position while straddling the moving conveyor belt.

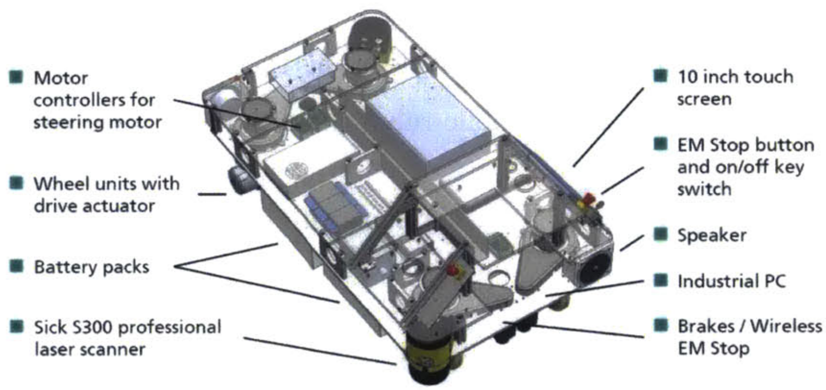

Figure 1-3: Schematic of the Rob@Work 3 mobile robotic base [11].

We select Rob@Work 3 mobile platform [11] as the basic system to build upon our robotic system for automotive assembly lines. Rob@Work 3 is a mobile robotic platform, developed by Fraunhofer IPA®, for industrial applications (see Fig. 1-3). Table 1.1 summarizes the key features of the Rob@Work 3 robotic platform.

Table 1.1: Rob@Work 3: Salient Features

Dimensions	$103 \times 57 \times 40$ cm
Weight	120 kg
Payload Capacity	150 kg
Maximum Speed	1.2 m/s
Actuators	4 wheels (2 motors per wheel, for driving and steering)
Sensors	Eight encoders (1 per motor) 2 SICK S300 Laser Scanners

The primary reason for the choice of this platform is the presence of four independently actuated wheels that can be both steered and driven. We leverage this characteristic of the platform for developing the control algorithm for robot motion on dynamic surfaces.

In addition the robotic platform has several features that are desirable for industrial applications, namely

- a high payload capacity that is required to carry out fetch-and-deliver and other tasks on the factory floor,
- a long battery life, that allows for extended robot operations without requiring frequent downtime for charging the batteries,
- on-board sensing, which includes motor encoders for odometry and two planar laser scanners with combined 360^o field of view,
- laser-scanner based safety system and emergency stops to prevent collisions with humans in the surrounding, and
- Robot Operating System (ROS) [54] based middleware, that enables easy adoption of off-the-shelf code for various robot functionalities, such as localization [28], mapping [33] and navigation [29, 49], and quick prototyping during system development.

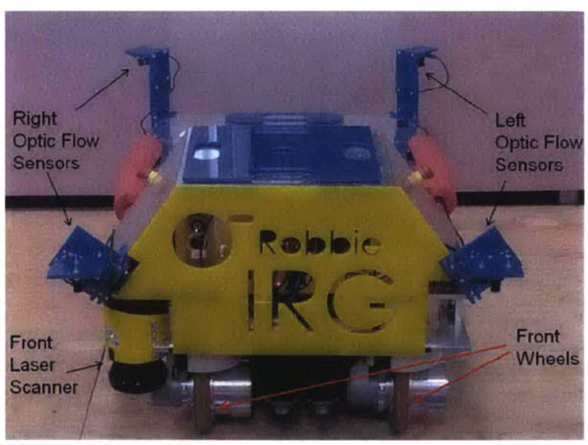

Figure 1-4: The modified Rob@Work 3 mobile robotic platform, with four optic flow sensors for augmented surface sensing capabilities.

Lastly, as described in Sec. 2.2, we augment the sensing capability of the robot mobile base by installing a suite of four PX4Flow optic flow sensors [39]. Figure 1-4 shows the basic robotic platform along with the augmented optic flow sensor suite.

1.2.2 Thesis in a Nutshell

In this thesis, we present a mobile robot designed for and capable of operating on automotive assembly lines along side human associates. To accomplish this aim, we developed a control strategy to enable autonomous navigation on dynamic surfaces, and design a sensing module capable of detecting the conveyor belts at sub-cm level accuracy. The developed control and sensing solution was implemented on the Rob@Work 3 hardware platform described in Sec. 1.2.1, and evaluated in an operational automotive factory floor.

In addition, with the aim of improving human-robot interaction in final assembly, we conducted a human subject experiment to assess the key factors impacting the saliency and collaborative fluency of a mobile robot assisting busy human co-workers. In this section, we provide a detailed outline of the thesis.

Trajectory Tracking on Dynamic Surfaces

We designed, implemented and tested a control algorithm and sensing system, described in Chapter 2, for the objective of tracking desired robot trajectories generated either by a teleoperator or path planner.

The sensing module senses the location and motion of dynamic surfaces using four on-board optic flow sensors and off-board contact-based wheel encoder (see Sec. 2.2). Various sensing alternatives for measuring the state of dynamic surface were considered, and first an on-board solution using only optic flow sensors was developed. To make the mobile robot more robust to the dynamic environment of final assembly, the sensing module was augmented with an off-board contact-based wheel encoder for sensing the speed of assembly line (dynamic surface). The knowledge of location and state of assembly line, the dynamic surface in automotive final assembly, is critical for position control of the robot on dynamic surfaces. In addition, the sensing information is useful for robot localization and for tracking objects, such as the cars being assembled, located on the assembly line.

Solutions readily exist for control of wheeled mobile robots on static surfaces; we build on the open-source Robot Operating System (ROS) software architecture of Rob@Work 3 [12] and generalize the control algorithm for environments with dynamic surfaces (see Sec. 2.1). The control algorithm utilizes the information from sensing module for tracking desired trajectories. Due to its modular design, the control algorithm can be used with any path planning algorithm irrespective of whether the robot is located on a static or dynamic surface. The developed trajectory tracking solution was first validated in simulation using the Gazebo robot simulator [44]. This was followed by hardware implementation of the control and sensing sub-systems. Initial hardware validation was carried out in a controlled lab environment, wherein a customized treadmill was used to emulate dynamic surfaces.

Evaluation in an Operational Automotive Factory

Next, to assess the performance of our robotic system in the dynamic and human-oriented environments, we evaluated the system in an operational automotive final assembly line (see Chapter 3). The trajectory tracking solution was integrated with algorithm for mapping, localization and path planning in dynamic environments. Evaluation in factory, among human co-workers, establishes preliminary success in the designed robotic system as evidenced through the quantitative metrics of navigational performance, and motivate future directions for improving the robot performance and interaction in time-critical tasks in human-centric environments.

Human-Robot Interaction: Collaborative Fluency and Robot Saliency

Fluent human-robot collaboration is critical to successful introduction of mobile robots in the automotive final assembly. Using the designed autonomous navigation capability for the robot, fetch-and-deliver will be one of the primary collaborative tasks that the robot can carry out. Since, our robotic system has a non-anthropomorphic geometry, omni-directioncal mobility and no in-built communication capability, we first aimed to evaluate its performance in a human-robot collaboration task to identify the need and format of design interventions.

We carried out a human-subject experiment, detailed in Chapter 4, which compared Human-Robot (HRI) and Human-Human Interaction (HHI) during delivery phase of fetch-and-deliver tasks. We observed statistically significant differences between human and the robot in the objective measures of fluency [37], as well as in a measure we defined as robot saliency. Interestingly, we observed that though the human-robot interaction was more salient it was less fluent than the human-human interaction. Additionally, participants were observed to respect a human collaborator's time more as compared to that of a robotic assistant. The study suggests a need for design interventions, especially for the time-critical operations in a factory where the robot idle time should be minimized and human-robot collaboration be made more fluent.

1.3 Related Work

There has been significant research towards designing and deploying interactive robots to assist humans. As exemplified in Fig. 1-5, these robotic systems span multiple domains. In this section, we provide an overview of such robotic systems with focus on robots designed for close proximity interaction with humans for the domain of final assembly.

1.3.1 Mobile, Interactive Robots

The breadth of recent and on-going work in robotics demonstrates a definitive push across sectors and academia to enable human-robot collaboration, including to enable novel applications, achieve ergonomic benefit, and improve task performance/efficiency. Indeed, several interactive robots haven been developed either as research prototypes and/or commercial products in the last two decades. Here, we describe few of these robotic systems. Please note that the intent of this description is not to provide an exhaustive list of such robots, but rather to ground our work and glean insights from the state of the art by studying existing, prototypical systems. We begin with a brief description of interactive robotic systems from domains other than manufacturing.

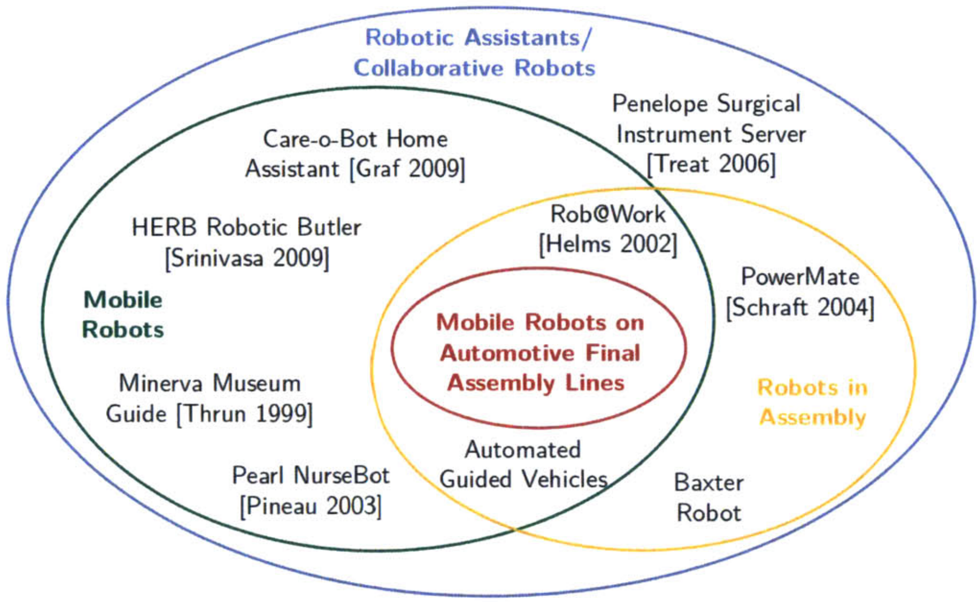

Figure 1-5: A Venn diagram exemplifying various interactive robotic systems. These robots span across multiple domains, including, health-care, homes, and final assembly. Further, many of these systems are mobile. However, our effort, to the best of our knowledge, is the first towards designing mobile robots that are capable of working along side humans on the automotive final assembly line.

Museums have been one of the first environments for deployment of interactive, mobile robots among humans. RHINO, an autonomous, interactive tour-guide robot, was successfully deployed in a densely populated museum for a period of six days as early as 1997 [52]. The autonomous operation of the robot was made possible due to various probabilistic algorithms, including those for localization, mapping, collision avoidance and planning. Minerva, the second generation of museum tour-guide robots, additionally included better planning and interactive capabilities [22]. These robots, developed and deployed more than a decade ago, demonstrate the potential of mobile robots for applications involving human-robot interaction.

More recently, interactive, mobile robots have been developed for applications in shopping malls [41], homes [7], health-care [15, 32, 53], and warehouses [34]. Along with being mobile, several robotic systems are capable of performing collaborative tasks with humans. For instance, the Home Exploring Robotic Butler (HERB) robot is capable of performing various manipulation tasks, including delivery of objects to humans [66].

1.3.2 Robots in Final Assembly

Industries have been the primary users of robotics at large scales. Indeed, today most robots are used in manufacturing operations [9]. As opposed to the robots described in Section 1.3.1, most of these industrial robots have been non-interactive. However, due to the advances that are enabling human-robot ineraction in service robotics, there also exists an increasing demand to incorporate mobile, interactive robots in the final assembly.

Indeed, to cater to the needs of modern factories, industrial robots are increasingly becoming more mobile [21], adaptive [13] and interactive [4]. Multiple robotic systems have been developed to collaborate with or assist humans in the final assembly [46]. Design of control strategies for mobile manipulation, with the aim of designing robotic assistants to aid human workers, is discussed in [42]. An autonomous mobile manipulator for assembly tasks is presented in [36]. The designed mobile system provides high flexibility and robustness, and has been empirically demonstrated to achieve high reliability in an insertion assembly task.

Robots in Automotive Final Assembly

Robotic systems have also been developed for our domain of interest, i.e., the final *automotive* assembly [56]. For instance, Akella et al. [16] discuss the design principles and prototype of cobots, an intelligent assist device (IAD) which provide ergonomic benefit to human workers in automotive assembly lines. Subsequently, these IADs have been commercialized for use in automotive assembly for assisting humans in tasks, such as, cockpit install and engine assembly. Apart from IADs, static robots have also been used for accomplishing specific tasks in car assembly, e.g., final door assembly [14]. Similarly, PowerMate, an workplace and time sharing robotic system, has been designed to assemble heavy parts of an automotive rear axle in collaboration with a human worker [57]. These robots operate in close proximity with humans; however, by being static are limited in the operations that they can perform and provide less flexibility while designing manufacturing processes.

Automated guided vehicles, as shown in Fig. 1-1b, are being used in many factories to carry out bulk of delivery of parts across the large-scale factory floors [10, 74]. However, these AGVs operate as line followers with their mobility being limited to the grids laid down in the environment, and are not capable of entering the automotive assembly line or performing tasks on moving parts/cars. The need and challenges of mobile robotic systems that can perform assembly tasks on moving assembly lines have been identified in [60]. As an enabling step towards developing these mobile robots, solutions for tracking moving objects from a mobile robot for applications in automotive assembly have also been developed [35, 61]. Using these tracking solutions and pure pursuit tracking control, Shi et al. have designed and implemented a mobile robot capable of performing assembly tasks on a moving vehicle body [59]. Specifically, the robot has been demonstrated to perform the task of attaching a wiring harness to a moving vehicle body.

The designed mobile robotic system, which is the closest work related to our design effort, is developed for applications in robotic automation stations [60] where the moving vehicle body is ferried by automated guided vehicles. Due to this the design assumes mobile manipulation being performed on static surfaces and does not consider the impact of dynamic surfaces on robot's operations. However, as described in Section 1.1, automotive final assembly lines typically include dynamic surfaces in the form of conveyor belts that ferry moving vehicle bodies. Further, the environment includes human associates that will work in close proximity with the robotic system, thereby necessitating consideration of safety and human-robot interaction. As detailed in subsequent chapters, in this thesis we consider both of these factors, i.e., dynamic surface and close proximity interaction with human associates, while designing and deploying a mobile robotic system for automotive final assembly lines.

Chapter 2

Trajectory Tracking on Dynamic Surfaces: Control and Sensing

Ensuring autonomous mobility is key to successful introduction of mobile robots for the final assembly of cars. The automotive final assembly presents a challenging environment which is dynamic, cluttered and largely unpredictable due to presence of moving cars and human agents in the environment (see Section 1.1 and Fig. 1-2). In addition, the automotive final assembly includes assembly lines, in the form of conveyor belts, which ferry cars around the factory floor. These assembly lines are essentially *dynamic surfaces* which could either be static or moving.

Autonomous mobility constitutes a robot being able to sense its environment, plan a path based on the sensed information to the desired goal location, and successfully follow the planned path. Solutions exist for each of these three components, i.e., sense, plan and control, when a robot is navigating a static surface [49, 63, 68]. However, no attempts have been made for achieving autonomous robot navigation on dynamic surfaces. In this chapter, we first present a modular control algorithm designed for trajectory tracking of Rob@Work-3 on a dynamic surface. The control algorithm allows the robot to successfully track a desired trajectory/path, generated either by a user or a path planner, irrespective of the surface being static or dynamic. As an input the designed control algorithm requires the location of the robot relative to the dynamic surface as well as the absolute velocity of the surface.

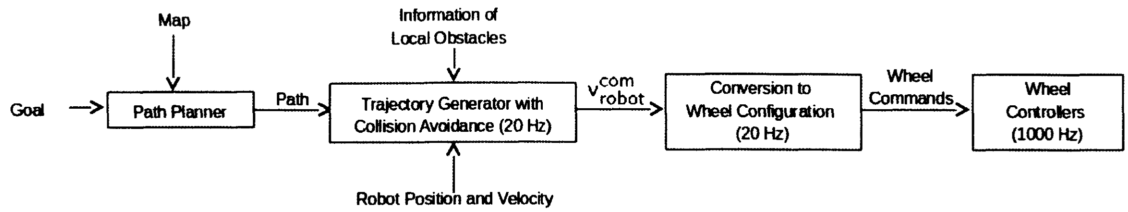

Figure 2-1: Nominal control architecture of the Rob@Work 3 robotic platform.

A number of exogenous factors influence the stopping and starting of the assembly line. Factories oriented around manual work do not typically maintain a reliable signal of the lines' *on/off* state and thus this state must be sensed directly by the robot. Thus, next a sensing subsystem is developed that provides location and speed measurements of the dynamic surface, an input for the designed control algorithm.

Other objects in the factory (part carts, cars, people) will move over different timescales and the robot will need localization and path planning algorithms that are suitable for dynamic environments. Due to the modular design of the control and sensing subsystems, the subsystems for dynamic surfaces can be coupled with a variety of (novel or existing) localization, mapping and planning algorithms thereby enabling autonomous mobility in the automotive final assembly line.

2.0.1 Reference Frames

The robot, due to its numerous parts and sensors, requires maintenance of over fifteen reference frames for its complete kinematic description. Primarily, three reference frames are used in this chapter: map, robot, and optic$_i$. The map frame, denoted by superscript M and used as reference for localization, is the static frame fixed to a reference point in the factory. The robot frame, denoted by superscript R, is fixed to the robot body, and is used for maintaining the robot odometry and location with respect to the static world frame. Finally, a reference frame is defined for each of the four optic flow sensors. This is used to suitably transform the sensor information into the robot frame.

2.1 Control on the Assembly Line

This section begins with a description of the nominal navigation system of our robotic platform. Though capable of successfully tracking a trajectory on a static surface, we observe through simulation that the robot is unable to do the same on surfaces that are dynamic. This motivates the need for a novel control subsystem, which is capable of navigating environments with both static and dynamic surfaces. We discuss the requirements and alternatives for designing such control algorithm, and present the design and performance of a modular control solution for our robotic platform.

2.1.1 Trajectory Tracking of Rob@Work 3

Control of wheeled robots is a well studied problem [63]. State-of-the-are wheeled robots are capable of following a given path or trajectory with high fidelity. The Rob@Work 3 robotic platform, too, can exhibit this trajectory tracking behavior [12] on static surfaces.

2.1.1.1 Nominal Architecture

Figure 2-1 shows the nominal control architecture of the robotic platform for navigation. The control architecture of the mobile base uses multiple feedback loops [24]. A path planning algorithm or a human tele-operator issues a desired trajectory, which is translated into velocity commands for the mobile robot. The desired velocity of the i$^{\text{th}}$ wheel $\mathbf{v}^{R}_{\text{wheel,i}} = (v_{\text{x,i}}, v_{\text{y,i}})$ is obtained in terms of robot velocity $(\dot{x}_{\text{r}}, \dot{y}_{\text{r}}, \dot{\phi}_{\text{r}})$ as follows,

$$v_{\text{x,i}} = \dot{x}_{\text{r}} - \dot{\phi}_{\text{r}} y_{\text{w,i}} \tag{2.1a}$$

$$v_{\text{y,i}} = \dot{y}_{\text{r}} + \dot{\phi}_{\text{r}} x_{\text{w,i}}. \tag{2.1b}$$

Each wheel velocity command is then converted to the wheel configuration, a steering angle ψ and angular rate $\dot{\theta}$ command, as detailed in [24]. A PD controller is used for controlling ψ and $\dot{\theta}$ for each wheel. The control architecture, implemented in ROS, also incorporates safety considerations that are essential while operating in an industrial environment.

2.1.1.2 Simulated Performance on Assembly Line

The control architecture of the Rob@Work 3 though capable of following a given trajectory is not designed for dynamic surfaces. However, the realized performance of the existing control architecture might still be acceptable due to the low speed (< 0.2 m/s) of dynamic surfaces encountered in the automotive final assembly. To test this hypothesis, we conduct simulation of trajectory tracking on dynamic surfaces using the nominal control architecture.

A test scenario, shown in Fig. 2-2, simulating a factory-like environment with a dynamic surface is created using the Gazebo simulator [44]. The dynamic surface, similar to assembly line in an automotive factory, either remains stationary or moves at a uniform speed in a fixed direction. Gazebo simulator is the natural choice for validation as it allows for testing of the hardware-ready ROS code in simulation.

In the test environment, the robot's task is to navigate across the dynamic surface that is moving with a constant velocity throughout the simulation. An off-the-shelf path planner [29] is used to generate the robot's path, which is a straight line in the current scenario.

Figure 2-3 shows four stills from the simulated performance of the nominal control algorithm while navigating the dynamic surface (a video of the complete simulation is available at `http://tiny.cc/aogo3w`). In spite of the low magnitude of surface velocity, the robot continuously deviates from the desired path while on the dynamic surface. As is evident from the simulated performance, the existing architecture fails to follow a desired trajectory on dynamic surfaces.

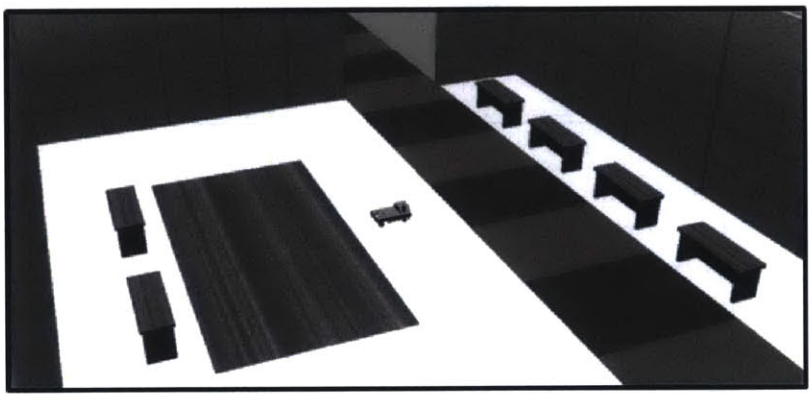

Figure 2-2: A Gazebo test scenario simulating factory floor with an assembly line.

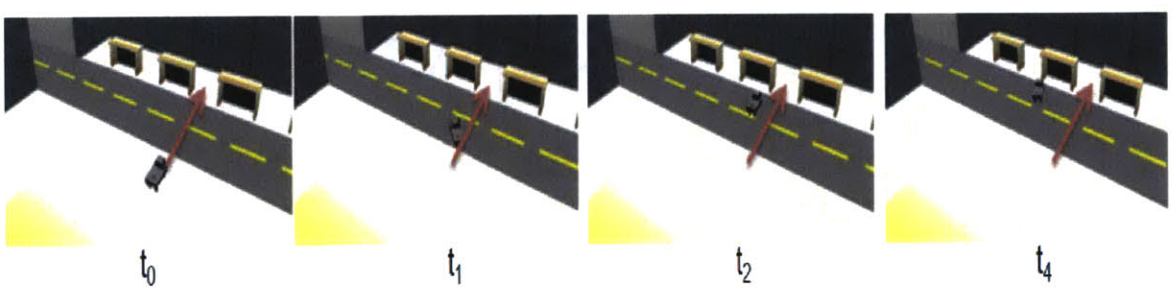

Figure 2-3: Four stills from the simulated performance of the nominal control architecture while the robot traverses the moving surface. The red arrow indicates the desired robot path. As the robot moves across the moving line, the deviation from the desired path continues to grow.

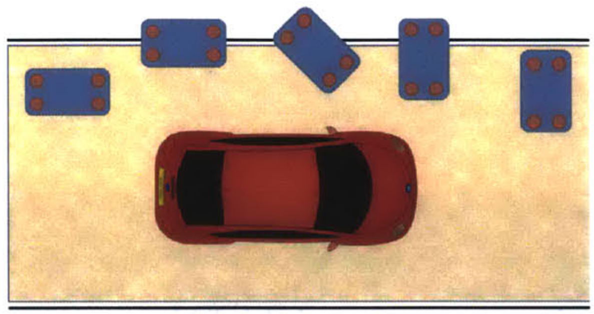

Figure 2-4: The robot can enter the conveyor belt in any arbitrary configuration. The figure shows few of the many possible orientations in which the robot (shown as a blue polygon with red wheels) can entering the conveyor belt that ferries cars. At any instant, any number of robot wheels can be present on the dynamic surface.

Effect on Robot Hardware

Not only does the nominal control architecture fails in tracking trajectories on dynamic surfaces, it also has a detrimental effect on the robot hardware. By not accounting for motion of the dynamic surface, the robot hardware experiences high torques while transitioning from static to dynamic surface or vice-versa. Repeated application of these high torques will structurally weaken the joints between the robot chassis and wheels, and may eventually cause the wheels to break off from the chassis. This impacts the maintainability of the robot and is highly undesirable for introduction of mobile robots on the factory floor.

2.1.2 Control Algorithm for Dynamic Surfaces

Having ascertained that the nominal control architecture is insufficient for robots navigating automotive final assembly lines, we aim to design a control algorithm which exhibits desired trajectory tracking performance on dynamic surfaces and avoids any undesired torques on robot hardware.

The mobile base can operate in multiple configuration with respect to the dynamic surface, i.e., it could have any number of its wheels on the dynamic surface in any arbitrary orientation (see Fig. 2-4). Further, due to its omni-directional motion the first wheel to encounter the dynamic surface will not be same for different trajectories.

Thus, along with functioning on the dynamic surface, the designed control algorithm should also be able to work when the robot is on a static surface, or is straddling the assembly line in any arbitrary orientation.

2.1.2.1 Control Alternatives

The nominal control architecture has been designed for static surfaces, and thereby does not model or account for effects of surface motion on motion or hardware of the robot. For successful trajectory tracking on dynamic surfaces, we consider control algorithms that explicitly account for effect of surface motion. Several approaches exist to model the effect of surface motion, either as an unknown but estimated parameter, as a disturbance, or as an sensed input.

High Bandwidth Position Control

To avoid any additional sensing requirement, one approach to achieve trajectory tracking is through a high bandwidth position control, i.e., to increase control gains to track desired positions along the trajectory. Though, such an approach is theoretically possible, through careful choice of control gains, it is highly undesirable from a hardware standpoint. A control which just corrects for position will theoretically be able to track the robot body center, but similar to the nominal control architecture will cause inordinate torques when the robot is straddling a dynamic surface. Such high torques may damage robot hardware, especially due to repeated robot motion on and off the dynamic surface.

Adaptive Control

An adaptive control approach [65] to trajectory tracking can also be applied by treating the surface speed as the unknown parameter. Design of independent adaptive controllers for each wheel will enable robot control even when robot is straddling the dynamic surface. However, while transitioning from static to dynamic surface or vice-versa any adaptive control design will exhibit transient response due to a jump in values of unknown parameter, i.e., speed of the surface. This, too, will result in repeated torques on the robot hardware which is undesirable on a factory floor.

Reference Shaping

To avoid any undesirable torques on robot hardware, one approach to trajectory tracking is to model the surface motion as an additional input to the system. Such a reference shaping approach would allow modular implementation, but require additional sensing of surface parameters. However, as this approach avoids any undesired effects on robot hardware, a key requirement for structural integrity of robots and introducing them on factory floor, we design the novel controller for trajectory tracking using this approach.

2.1.2.2 Controller Design

We design a control algorithm, based on reference shaping, which considers the speed of dynamic surface as an additional input. Fig. 2-5 shows the designed control architecture, and Algorithm 1 details the modification to the nominal controller. This control architecture leverages the independent actuation of each wheel, and compensates for motion of the dynamic surface through suitably modifying the reference to the robot's wheel controllers. This modification results in a modular design that preserves use of the existing wheel PD controllers, and software architecture.

Algorithm 1: Modification to the command to nominal wheel controller for navigating environments with dynamic surfaces, such as, automotive final assembly lines

Input: Nominal command to the wheel controller
Output: Compensated command to the wheel controller
$\mathbf{v}^{\text{robot}}_{\text{surf,i}}$: absolute velocity of the surface at i^{th} wheel in robot frame
$\mathbf{v}^{\text{robot}}_{\text{wheel,i}}$: absolute velocity of the at i^{th} wheel
foreach *robot wheel* **do**
  sense absolute surface velocity ($\mathbf{v}^{R}_{\text{surf,i}}$) at wheel;
  modify the nominal command:-
  $\mathbf{v}^{R}_{\text{wheel,i}} = \mathbf{v}^{R}_{\text{wheel,i}} - \mathbf{v}^{R}_{\text{surf,i}}$;
end

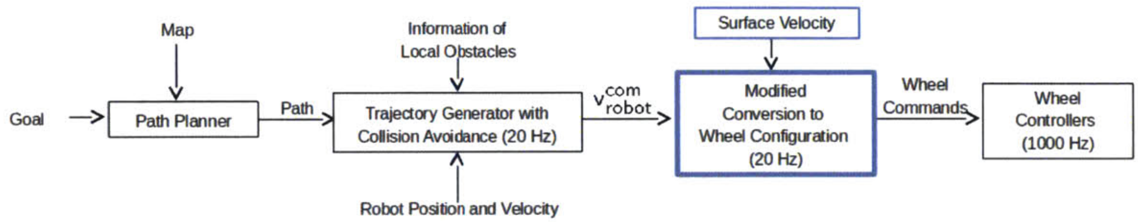

Figure 2-5: Designed control architecture of Rob@Work 3 for environments with dynamic surfaces.

The nominal wheel controllers, described in Sec. 2.1.1.1, operate assuming the commands are issued for the case of static surface. Hence, if any of the wheels are on a moving surface, the nominal controller does not provide the desired response. To overcome this issue, we compensate the command for each wheel ($\mathbf{v}^R_{\text{wheel,i}}$) based on the absolute surface velocity at its point of contact ($\mathbf{v}^R_{\text{surf,i}}$). Algorithm 1 describes how the commanded velocity for each wheel is altered assuming the knowledge of surface velocity. The modified wheel velocity command is used to compute the desired wheel configuration (ψ,$\dot{\theta}$).

2.1.2.3 Simulated Performance on Assembly Line

To validate the modifications to the control algorithm prior to its implementation on the robot hardware, we use the test environment developed in Section 2.1.1.2 using the Gazebo simulator. To validate the controller independently of the sensing, the Gazebo simulation assumes perfect sensing of the surface velocity. Robot's task and environment are kept identical to Section 2.1.1.2.

Fig. 2-6 shows the deviation of robot's center from the desired path for the simulated scenario. Using the designed control subsystem, the maximum deviation is observed to be <4cm for this task when the moving surface is operating at 0.10 m/s (a representative value for automotive assembly lines). Fig. 2-7 shows four stills from the simulated performance of the designed subsystem (video available at `http://tiny.cc/aogo3w`). In contrast to the nominal control architecture's simulated performance shown in Fig. 2-3, using the modified system the robot can successfully navigate across the simulated assembly line. This is accomplished by

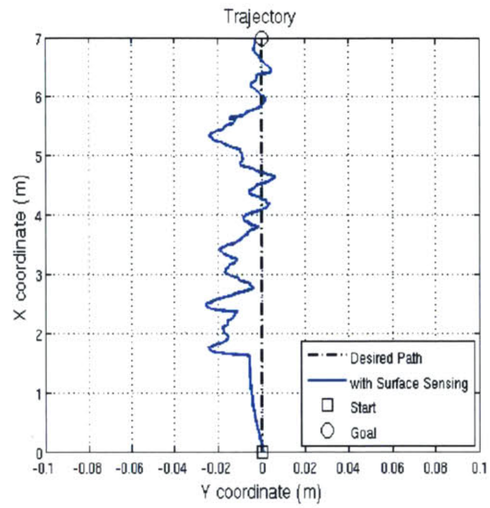

Figure 2-6: Performance of the modified system for the task of robot traversing the moving surface from a typical Gazebo simulation run. Magnified view shows the deviation in the path as the robot enters the moving surface. However, this deviation is <4cm for the entire path, validating the modification in control sub-system.

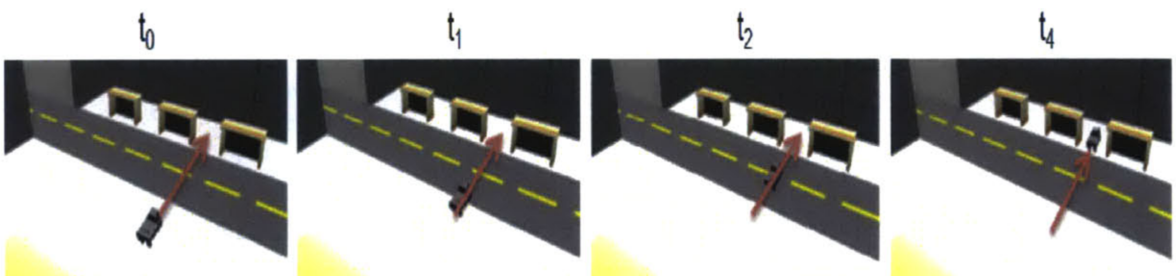

Figure 2-7: Four stills from the simulated performance of the modified control architecture while the robot traverses the moving surface. The red arrow indicates the desired robot path. As the robot moves across the moving line, the deviation remains bounded with the designed control architecture.

dynamically compensating for the surface velocity and correcting the robot heading. Next, we describe the design of the sensing subsystem that provides the information regarding the dynamic surface, required by the designed control architecture.

2.2 Sensing the Assembly Line

Awareness of the dynamic surface and its current state is an important requirement that a mobile robot needs to fulfill to operate on automotive final assembly lines. Primarily, for executing the control algorithm described in Section 2.1, the robot needs to know whether it lies on the assembly line or not. Additionally, for accomplishing any task on the moving line, such as, assisting in car assembly, the robot needs to maintain an estimate of the state of the assembly line.

Requirements: The key requirements for designing a sensing subsystem include maintaining an online estimate of assembly line location, and measuring the absolute speed of assembly line. As maintaining these estimates are critical to robot's function, safety capability and structural health it is important to have a reliable and robust sensing module. Lastly, to minimize the need of any additional infrastructure, it is desirable that the sensing be done on-board as far as possible.

In this section, we explore different sensing alternatives, design an on-board sensing system and augment it with off-board sensing to improve the reliability of the sensing module.

2.2.1 Sensing Alternatives

We explored the use of four types of on-board sensors for measuring the speed of the dynamic surface: miniature radars [6], optic flow sensors [39], contact-based encoders, and inertial sensors. As the surface moves relatively slowly (< 0.2 m/s), the performance of miniature radars and low-cost inertial sensors is limited by accuracies at low speeds. Further, measurements from an indirect method (such as, inertial sensor based system) will be reactive, detecting surface motion only via disturbance in the robot's motion caused by the surface. This will result in a delayed response of the controller. An on-board contact sensor will negatively affect the ground clearance of the robot, limiting its mobility over uneven surfaces.

On-board optic flow sensors have been previously used for maintaining location estimates of mobile robots [50] on static surfaces. Here, instead, we intend to apply them for detecting the surface velocity. Further, the images from on-board optic flow sensors can also be potentially used to detect the location of robot relative to the automotive final assembly line. In our design, we first evaluate the use of on-board optic flow sensors for sensing the assembly line.

An on-board optic flow sensor can only provide information regarding the relative speed of the dynamic surface relative. As detailed in the following sections, the optic flow sensors will need to rely on an independent source of robot velocity to measure the absolute surface speed. Though, an estimate of robot velocity can be maintained using other sensors, e.g., on-board laser scanners, a more robust solution is desired for application on the factory floor. Hence, to increase the sensing reliability, we include an off-board contact-based wheel encoder in the final design of our sensing module.

2.2.2 Location of Assembly Line

The initial estimate of the location of the assembly line is known to the robot *a priori*, based on the static map of the environment. Given that the location of the assembly line is fixed, even though the surface is dynamic, ideally this a priori information should be sufficient to maintain an estimate of assembly line location. For instance in Fig. 1-2, the surface of conveyor belt (which is wooden in texture) which ferries cars moves, but the boundary between the conveyor belt and the static surface (gray in color) does not change.

However, the assembly line location is to be known relative to the robot, which in turn requires knowledge of robot's pose and map representation. This estimate of robot's pose is obtained using an online localization algorithm, which is informed by noisy sensors and thus has some estimation error and drift. Thus, the accuracy of the a priori estimate of the relative location of the assembly line with respect to the robot will only be as good as localization accuracy. The localization algorithm is designed and implemented with the specification of maximum localization error as 5cm.

Figure 2-8: PX4Flow Optic Flow Sensors [39]

As reliable and accurate estimation of assembly line location relative to the robot is important for maintainability of the robot hardware (see Section 2.1), we need to maintain an estimate of the assembly line location in the current map representation of the robot which is updated online. To obtain this online estimate, we use four PX4Flow optic flow sensors [39] mounted on-board the robot and facing downwards.

The PX4Flow, shown in Fig. 2-8, is a CMOS image based optical sensor designed for mobile robotic applications, including aerial vehicles [39]. The sensor includes a microcontroller, image sensor, gyroscope and sonar with on-board processing, and interfaces with the robot through ROS[1]. Presence of four sensors allows for detection of assembly line independent of the robot's heading or pose while it enters the assembly line; this is especially of importance due to the omni-directional motion of the Rob@Work 3 mobile base (see Fig. 2-4).

These optic flow sensors are integrated with the robot software using ROS and transmit black-and-white images of the surface to the robot computer with a frequency of 6Hz. We use standard image processing techniques to detect if the image includes a line corresponding to the boundary of the assembly line. Note that such a line will be present in the image when the sensor transitions from static to dynamic surface or vice-versa. Specifically, for each image, transmitted by the four sensors, we use Canny edge detector[2] to detect edges in the image. Next, we calculate the Hough transform [3] to detect lines in the detected edges. Lines which differ from the expected orientation of the assembly line are eliminated from the possible detection of the assembly line.

[1]`http://wiki.ros.org/px4flow_node`
[2]`http://goo.gl/AOFVvV`
[3]`http://goo.gl/o10ssU`

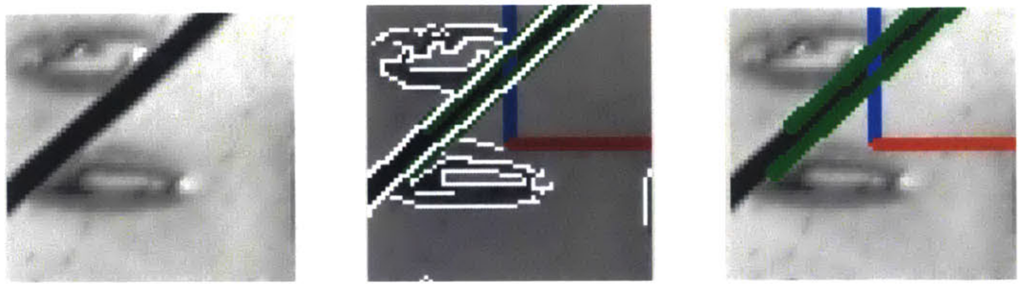

(a) Raw image obtained from the optic flow sensor

(b) Edge detection using Canny edge detector

(c) Line detection in the image using Hough transform

Figure 2-9: Image processing pipeline to detect the location of assembly line

Lastly, based on the current location of the sensor and the line in the image, the location of the assembly line is updated. This process mitigates the error in assembly line estimation based purely on localization information, achieves sub-cm level accuracy, and provides redundancy in detection of assembly line during robot operations.

2.2.3 Speed of Assembly Line

2.2.3.1 Optic Flow Sensors

Along with detecting the location of assembly line in robot's map, the image sensors of the PX4Flow also provide the surface velocity. In the off-the-shelf PX4Flow sensor module, the image gradient, obtained from the CMOS camera, is processed using a Kalman filter implemented on the microcontroller to obtain a refined estimate. The derived optic flow is next compensated for image distance (using sonar), and sensor rotation (using on-board gyroscope) to provide the image velocity.

In our application, by being mounted facing downwards, the sensors capture images of the surface underneath the sensor, and provide its velocity relative to that of the robot. Use of optic flow sensors requires that the surface have sufficient image features, and we empirically validate that this is the case for our particular factory environments. The sensors are mounted at a fixed height and maintain a constant distance from the surface. Hence, the fixed height of the sensor is used in calculating velocity, rather than the noisy sonar measurements. This modification considerably reduces the noise in the velocity measurement for our application. We then process the output velocity using a low pass filter to generate the estimate of surface velocity.

The sensors are attached to the robot frame to measure surface velocity relative to the robot at the ith wheel ($\mathbf{v}^{\text{optic}_i}_{\text{surf}_i\text{-robot}}$), which is then transformed from sensor frame, optic$_i$, to the robot frame, R, to provide $\mathbf{v}^{R}_{\text{surf}_i\text{-robot}}$,

$$\mathbf{v}^{R}_{\text{sensor}_i} \triangleq \mathbf{v}^{R}_{\text{surf}_i\text{-robot}} = \mathbf{v}^{R}_{\text{surf}_i} - \mathbf{v}^{R}_{\text{robot}}$$

Robot localization using only laser scanner measurements is a problem that has been studied in detail [17, 45]. Successful algorithms have been tested in real systems and have been developed and implemented. For our application, we use a scan matching approach to localization as described in [45], primarily because of its open source ROS implementation[4]. Further, the algorithm requires low computational resources and is capable of simultaneous localization and mapping. Implementation of the algorithm on our robotic hardware results in acceptable localization performance in static environments. Robot's absolute velocity, obtained from the laser-scanner based localization algorithms, is then combined with the sensor measurement to calculate the absolute surface velocity,

$$\therefore \mathbf{v}^{R}_{\text{surf}_i} = \mathbf{v}^{R}_{\text{sensor}_i} + \mathbf{v}^{R}_{\text{robot}}.$$

Performance

The performance of the PX4Flow sensor is first compared to the ground truth information, to gauge its accuracy and inform the requirement of additional processing of the raw measurements. Fig. 2-10 shows the performance of the optic flow sensor for a challenging practical scenario: where the moving surface starts from rest and then reaches its maximum velocity, and then again resumes its original rest state.

[4] http://wiki.ros.org/hector_slam

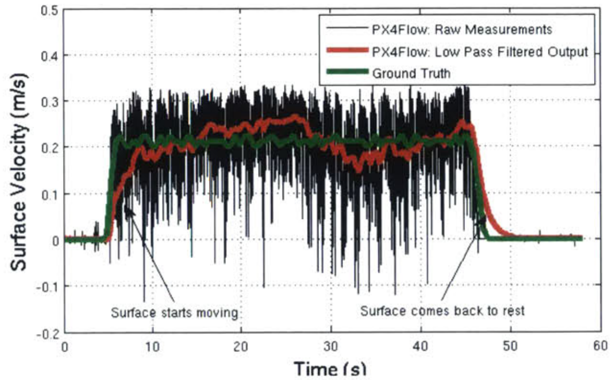

Figure 2-10: Performance of the optic flow sensor as compared to the ground truth. Both raw measurements from the optic flow sensor, and its filtered value using low pass filter are shown.

As can be observed from the Fig. 2-10, the sensor is capable of tracking the surface motion albeit with substantial noise. Further we observe that performance and measurement range of the optic flow sensor is dependent on the quality of image features present on the surface.

High levels of noise in surface sensing will produce cascading effects in the controller performance, and rapidly changing commands will negatively affect the actuator hardware. This motivates the design of a discrete time Kalman filter with its state as the surface velocity. The process model accounts for the variation in the state through a discrete white noise term with variance σ^2. Process noise is characterized based on the variation in the surface speed during its on state, which for our application is approximated as $3\sigma = 1$ cm/s. The measurement noise is also assumed to be white for filter design, and its standard deviation is approximated as 3 cm/s, based on the in-house tests. A Kalman filter with steady-state gains is used rather than one with time varying gains. This is equivalent to a low pass filter with gains set to the steady-state gains of the Kalman filter.

The performance of the optic-flow based surface speed detection, though acceptable for static test environments, is directly dependent on the localization accuracy of the laser-scanner based localization algorithm. In dynamic environments accuracy of such localization algorithm can degrade, resulting in a noisy estimate of surface speed from the optic flow sensors.

2.2.3.2 Off-board Contact-based Sensor

To overcome the limitation of optic flow based sensing in dynamic factory environments, in our final design we use an off-board wheel encoder mounted on the dynamic surface which measures the speed of the assembly line. The sensing information is transmitted wirelessly to the robot base, and provides a robust solution to detect the speed of the assembly line without the need of any additional measurement of robot velocity.

Though presence of an off-board sensor requires additional infrastructure such as wireless communication, the need for system to be robust far outweighs the limitation of these additional infrastructure. This simple solution also alleviates any cascading effects of error in estimating robot velocity which are imperatively present in any on-board solution for estimating the speed of dynamic surface.

2.2.4 Initial Hardware Validation

After independent validation of the controller (through Gazebo simulation) and optic flow sensor (via comparison with the ground truth), we proceed towards observing the performance of the combined system. We use a customized treadmill[5] to test the system on a real moving surface. The velocity of the treadmill can be controlled, thereby enabling testing at different operating conditions.

As an initial step, the sensor for surface velocity is mounted off-board and performance of the system for the 'position hold' task is observed. Position hold is an important task for a robotic assistant, especially while delivering parts to a human associate working on a moving line. Fig. 2-11 shows the two test scenarios for the position hold experiment, namely, when the robot is completely on, and when it straddles the moving surface. For the position hold tests, the customized treadmill operates at $\approx$20cm/s during its on state. For both the cases, the robot successfully maintains a constant position despite the change in the assembly line's on/off state.

[5] `http://goo.gl/8JGtVQ`

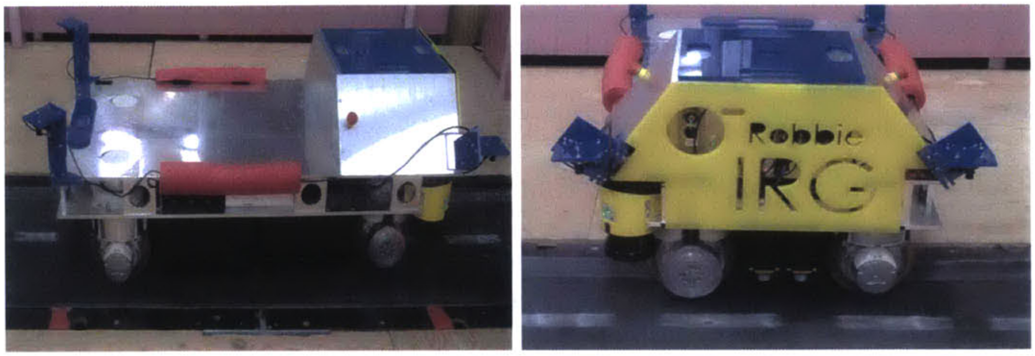

Figure 2-11: Position hold experiments with off-board surface sensing being conducted in two cases: the robot being completely on (left), and straddling (right) the moving surface.

In these tests, the laser scanner based localization is used to detect when the wheels are on the customized treadmill and the surface velocity is obtained using the ground truth sensor; for surfaces with sufficient image features off-board optic flow sensors also yield similar performance. Next, the robot's ability to navigate across the moving assembly line is tested in the same test environment. Similar to Gazebo simulation, the robot successfully traverses the moving surface (accompanying video available at `http://tiny.cc/aogo3w`).

2.2.5 Summary and Next Steps

In this chapter, we presented a novel control and sensing solution for a mobile robot to track desired trajectories on dynamic surfaces, such as those encountered on automotive final assembly lines. Information about surface velocity is essential for our modified control algorithm. Off-the-shelf sensors and localization algorithms are explored for fulfilling this requirement, and a customized solution is presented using PX4Flow optic flow sensors and an off-board contact-based wheel encoder.

We perform initial validation of the controls and sensing subsystems using both software simulation and hardware tests. Software simulation and hardware implementation of the full hardware system yield promising results; using off-board surface sensing the robot successfully moves across and maintains its position while on, and while straddling, the moving line.

In the current test setup, we demonstrated operation of the robot on a moving surface but in an otherwise static, known environment. In a real factory setting, the environment will be dynamic with part carts and humans moving in a very dense and cluttered space. Additional capabilities are necessary to detect and avoid humans and other obstacles, and to ensure efficient robot navigation. The following chapter discusses these challenges associated with deployment of our robotic system in an operational factory floor.

Chapter 3

Robot Evaluation on the Automotive Final Assembly Line

Having developed the control and sensing system that enables robot navigation on dynamic surfaces and testing the same in controlled laboratory environments (as described in Chapter 2), here we focus on deploying and evaluating our mobile robotic system in an operational automotive final assembly. Several challenges need to be overcome when the robot enters the dynamic, human-centric factory environment; these span across subsystems including localization, mapping, sensing, path planning, and control as well as their interactions.

This chapter discusses these challenges and presents the results from a five-day long deployment of the mobile robotic system in an operational factory floor. We begin with a description of the test scenario that was used for evaluation of the robot on the automotive final assemble line. Next, salient details of various subsystems that were instrumental in deploying the robot are discussed. Results from the factory demonstration are presented next, which include quantitative analysis of the navigation performance of the robot on automotive assembly line as well as qualitative observations of close-proximity, human-robot interaction. These demonstrations establish confidence in the designed system and motivate future directions for improving the robot performance and interaction in time-critical tasks in human-centric environments.

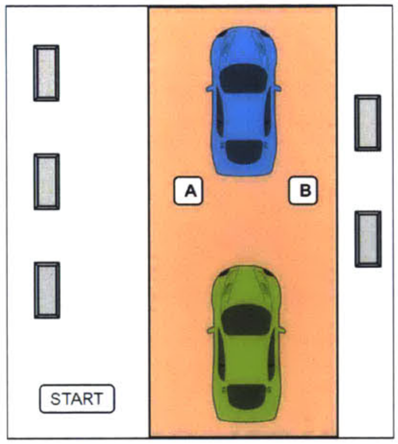

Figure 3-1: A schematic of the robot's work environment during the factory evaluation. The work environment was shared with human associates working on the assembly of cars. The figure depicts important landmarks - start, A, B - associated with the robot's task. Locations A and B are defined relative to the moving car, and are dynamic points in the absolute frame. The dynamic surface is shown in orange.

3.1 Test Scenario

As the next step in our evaluation and to subsequently guide our design, we deployed the robot in an operational automotive final assembly line. This section describes the details of the specific test scenario used to demonstrate and evaluate the designed mobile robotic system. The test was carried out over a period of five days. Though name and location of the factory cannot be disclosed to maintain anonymity, we provide the features of the test scenario pertinent to the system demonstration.

Environment of the Robot

The test was performed in a work area, where each car roughly spent between 150-180 seconds (i.e., the cycle time) due its motion on the dynamic surface. A schematic of the robot's work area in the factory floor is depicted in Fig. 3-1. The motion of the assembly line was beyond the control of the robot and was dictated by the schedule of the factory. The line typically moved continuously at an average speed of approximately 8 cm/s and stopped during certain hours according to the factory's schedule. Additionally, we also observed stops of short duration due to unplanned/unscheduled events in the factory. As the motion of the dynamic surface was not modified to suit the robot demonstration, the robot had to adapt to the factory's schedule as would be expected of an autonomous system in a typical work environment.

Human Co-workers

The work area included human associates usually working on the assembly of cars adjacent to that of the robot and at times simultaneously on the same car. The human associates working along side the robot on assembly of car were briefed about the presence of the robot on the first day of the task. They were asked to work naturally along side the robot, except not to interfere with the robot motion when it was moving unless necessary.

Navigation Task

The robot's task involved going from the start location on the static surface, i.e., outside the conveyor belt, to location A next to an on-coming car. Location A is defined relative to the car and moves in the absolute frame requiring the robot to track and plan path to a moving goal. After reaching at the location A, the robot had to perform a task for car assembly and return back to the static surface.

Details of the specific assembly task are proprietary, but from perspective of navigation the task required the robot to perform position hold relative to the moving car while the task was being performed. This manufacturing task had to be completed within the cycle time of the car being in robot's work cell, making the task highly time-critical. As a more challenging task, we also evaluated the robot performance when it had to perform the same task on either side of the car within the single time. The robot motion in this task involved the following waypoints: start, A, B, A, start.

During the test at least three engineers were present to monitor the robot and override the autonomous robot in case of any unwarranted behavior. The task was initiated by one of the engineers following which the complete task was performed autonomously. Next, we describe the different robot subsystems designed, implemented and integrated to evaluate the robot in the operational factory floor.

3.2 Enhancements and Refinements

As detailed in Chapter 2, we developed enabling solutions for robot navigation on dynamic surfaces. Prior to the factory deployment, these developed solutions were refined and improved through both software simulation using Gazebo simulator [44] and hardware validation in static, controlled environments.

In this section, we describe the refinements and enhancements to the developed control and sensing solutions as well as the supporting subsystems (mapping, localization, path planning, etc.) while deploying the robot in an operational factory. Though successful during the preliminary validation in static environments, these solutions were revisited with the aim of making them more robust for operations in dynamic environments among human associates and expensive factory equipment.

3.2.1 Trajectory Tracking

Due to the modular design of the control algorithm, no modifications are required to the motion control architecture for use in the dynamic factory environments. Thus, for the test scenario we use the developed control architecture as is when the robot navigated to and from the assembly line.

To improve the robot's position hold performance while it is performing the assembly task, however, the motion commands are issued relative to the car's reference frame (as opposed to the absolute reference frame) once the robot reaches location A/B. This essentially results in application of no control effort when the assembly task is being performed, with better position hold performance and identical control architecture.

For the sensing module, parameters of the image processing algorithm (described in Section 2.2.2) were tuned based on the surface properties of the factory environment. This helped improve the the location sensing of assembly line. The overall architecture of the sensing subsystem, similar to the control sub-system, remained identical.

The contact based sensor, i.e., the off-board wheel encoder, to sense the surface speed was placed on the assembly line next to robot's work cell and connected to the on-board robot through a wireless network.

3.2.2 Supporting Subsystems

Along with the control and sensing solutions developed in Chapter 2, the preliminary validation required use of off-the-shelf algorithms for robot localization, mapping and path planning. These off-the-shelf solutions though useful for initial validation, were substituted by validated, proprietary software developed by Fraunhofer IPA [1] for robot deployment in dynamic environment.

Additionally, in contrast to the preliminary validation, the factory task required the robot to track and plan paths to moving objects, i.e., the cars being assembled. This required use of algorithms, also described subsequently, to track moving cars and plan paths to moving goals for the factory task. Here, to provide a complete picture of the developed system, we provide an overview these subsystems including their input and output information.

Localization

The localization algorithm requires knowledge of the static map of the environment and provides on-line update of the robot position and orientation every 0.01 s. For our factory deployment, a map of the environment was created while the assembly line was static and was used as an input for the localization algorithm. It uses sensing information available from the on-board laser scanners, encoders on robot wheels, and the sensing module developed in Section 2.2.

Encoders on the robot wheels are used to provide odometry information, however this information can be ambiguous on the dynamic surface due to robot exhibit motion in the global frame (due to the dynamic surface) in spite of the wheels being stationary.

$$\mathbf{v}^R_{\text{odom,rectified}} = \mathbf{v}^R_{\text{odom}} + \mathbf{v}^R_{\text{surf}} \tag{3.1}$$

[1] http://www.care-o-bot.de/en/publications.html

Hence, the algorithm additionally uses information regarding the speed of assembly line, as measured in Section 2.2, to compensate for robot motion due to dynamic surface and obtain disambiguated robot odometry ($\mathbf{v}^R_{\text{odom,rectified}}$). Equation 3.1 describes how surface velocity ($\mathbf{v}^R_{\text{surf}}$) is used to compensate the robot's odometry information ($\mathbf{v}^R_{\text{odom}}$), and thereby alleviate the error in robot's odometry caused due to the moving surface.

Though the localization algorithm takes static map as an input, it is applied to provide localization estimates in dynamic environment which includes humans, cars and other dynamic objects. Due to sensing inaccuracies, dynamic surface and dynamic environment the localization error can grow as high as 5cm during typical operations.

Path Planning

For path planning, an implementation based on the Elastic Band planner [55] is used. This reactive planner requires knowledge of the robot goal and current obstacles in the environment, and provides smooth paths that avoid obstacles in the current environment. Though the path planner does not reason about future states of the environment while creating robot paths, it is used for its ability to quickly replan and create smooth robot paths. These properties are desirable in the dynamic, human-centric environment of the factory floor. In addition, the planner is augmented with the capability to plan paths to dynamic goals, such as, the moving cars on the assembly line.

Car Tracking

Creating path plans to the dynamic cars located on the assembly lines requires the robot to be able to track these cars. Hence, a tracking algorithm is implemented to track cars on the assembly line. The tracking algorithm first uses laser scanners to identify the cars that are to be assembled, and uses a filter to track the identified cars based on the information received from laser scanner and speed sensing of the dynamic surface.

3.3 System Performance on the Automotive Final Assembly Line

In this section, we describe the performance of the designed robotic system from the final day of the five day long robot deployment in the operational automotive final assembly line. We first describe the performance of individual subsystems, followed by the navigation performance of the overall robotic system. Lastly, since, the robot operation were carried out in an operational factory we report the observed interaction between human associates and robot.

3.3.1 Task Characteristics

The presented results are based on data logs from fifteen instances of the test scenario, described in Sec. 3.1, recorded during the factory evaluations. Figure 3-2 shows the task times for each of these runs. As can be seen from the figure, the task time for trial # 9 was significantly higher than others; this trial of the task depicts the more challenging scenario when the robot had to perform assembly tasks on either side of the car, i.e., at both locations A and B. On average for these trial runs the single-side task took 53.77 sec, as required due to the finite length of robot's work cell and cycle time of each car within it.

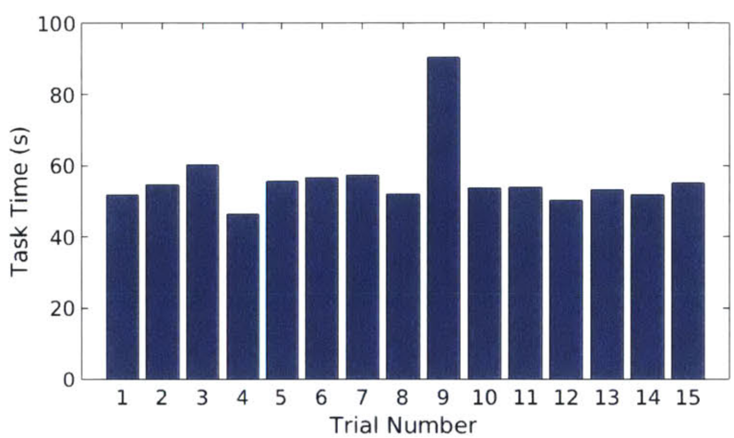

Figure 3-2: Task times across different trials of robot's test scenario.

Delay in this process might result in some cars not being assembled, thereby causing delays in the manufacturing process and affecting task efficiency and manufacturing quality. The ability of the robot to navigate within the cycle time of 150-180 seconds allows for robot to perform tasks for the final assembly of cars.

3.3.2 Assembly Line: Location

In order to use the designed control and tracking algorithms for of the dynamic surface, the robot needs to know which robot wheels are on the dynamic surface. A typical localization algorithm includes error arising from two sources : dead reckoning and odometry drift [2]. By periodically updating the location of the transition from static to dynamic surface, we can significantly reduce the effect of localization drift error in detecting whether the robot is on the assembly line or not.

For the demonstration of the robot in the factor floor, Figure 3-3 shows the number of assembly line location updates in each of the trial runs, using the optic flow sensors and image processing pipeline described in Section 2.2. Typically, the robot detects this transition at least once during each task, thereby allowing for better detection of which wheels are on the assembly line. In a few runs we also observe the transition was not captured; this is because of the limited update of the images from the optic flow sensors (6 Hz), and can be rectified by pubilshing surface images at a higher rate.

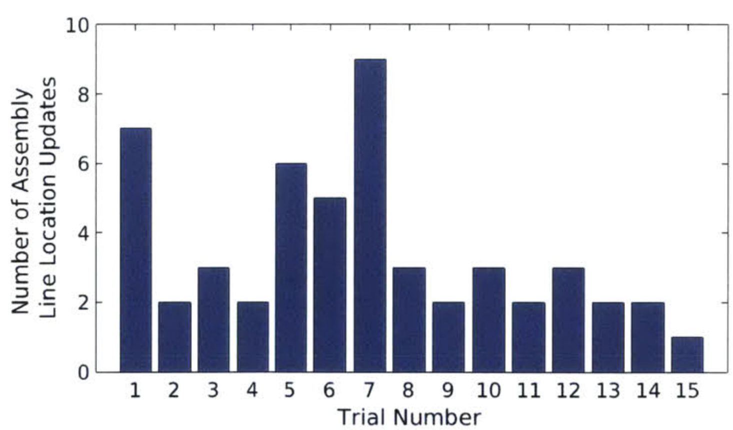

Figure 3-3: Number of assembly line location updates using the optic flow sensors across different trials of robot's test scenario.

3.3.3 Assembly Line: Speed

Figure 3-4 shows the typical speed profile of the assembly line as it transitions from off to on state. As, no ground truth information is available from any other sensor, we report only the measured information from the wheel encoder. The typical, measured speed of the assembly line when in on state was 7.89 cm/s with a standard deviation of 0.17 cm/s. The speed information of the assembly line was available to the robot throughout its motion at an update rate of 30 Hz.

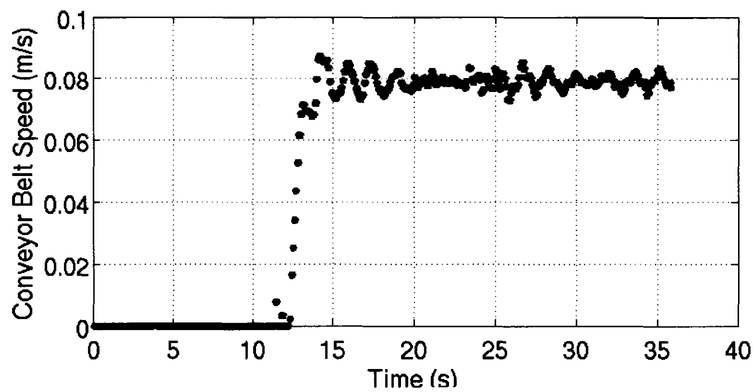

Figure 3-4: Typical speed profile of the assembly line.

3.3.4 Control Subsystem

To mitigate adverse torques on the robot hardware, the algorithm modifies commands to the robot wheels based on the location and speed measurements of assembly line. As a fool-proof check for the control subsystem we tested the robot in a position hold task when it was straddling the dynamic surface similar to Section 2.2.4. This scenario is specifically challenging due to the excessive, damaging torques that the robot could experience in the case the control algorithm does not compensate for the surface motion. The robot was observed to successfully hold its position despite only two of its wheels being on the moving conveyor belt.

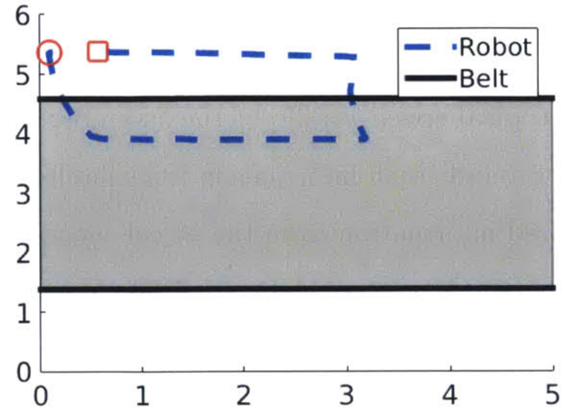

Figure 3-5: A typical robot trajectory for the test scenario with assembly task being performed only on the left side of the moving car (location A). The gray area denotes the dynamic surface. The red circle denotes the start location, and the red square denotes the final location of the robot. Once the robot reaches next to the car, it remains stationary relative to the car. This is observed as the straight line motion of the robot on the assembly line.

3.3.5 Navigation Performance

Having discussed operation of the individual subsystems, we present the overall navigation performance of the integrated system. Figure 3-5 shows the robot path from one run of the robot tasks on the automotive final assembly line. Similarly, Fig. 3-6 shows a similar plot when the robot performed the assembly task on either side of the moving vehicle. Further, during certain trials we observed that the dynamic surface changes it state during the task; for instance, the assembly line (and consequently the car) stops moving after the robot begins its motion. By maintaining an estimate of the speed of the assembly line, the robot can accomplish the task despite such unscheduled changes in the state of the assembly line. In all the recorded runs of the designed system, the robot was successfully able to navigate on the dynamic surface of the assembly line irrespective of whether it was moving or stationary.

An accompanying video of the rviz visualization of the mobile robot for both of these tasks is available at `http://tiny.cc/c94gyx`. Due to proprietary limitations, we show the rviz [2] visualization of the test environment and the robot's task instead

[2]http://wiki.ros.org/rviz

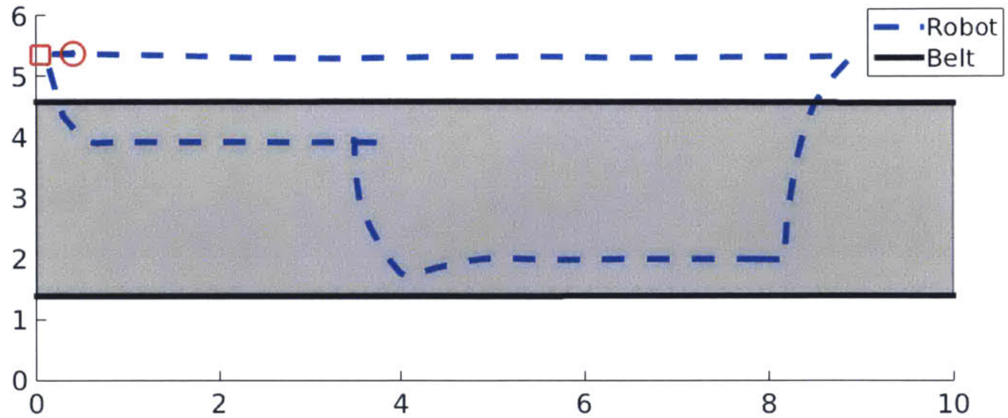

Figure 3-6: A typical robot trajectory for the test scenario with assembly task being performed on both sides of the moving car (location A and B). This task requires the robot to travel farther and spend more time on the dynamic surface.

of the raw video. This visualization has been created using the laser scanner and robot odometry information recorded during the evaluation.

3.3.6 Qualitative Observations: Human-Robot Interaction

An interesting by-product of evaluating our mobile robot in an operating factory floor was to observe the naturalistic interaction between the factory associates and the mobile platform, which is rather challenging to replicate and achieve in controlled laboratory experiments. The human associates working in close proximity to the robot were briefed about the robot's operation on the assembly line, and that its motion was being monitored by the on-site engineers.

We observed the human associates working next to the robot though showed some initial curiosity as and when the robot performed any novel sub-task, their behavior eventually became indifferent to the robot's motion. Over the period of the work week, during the evaluation period, their deviation in behavior diminished as the robot began completing its task more reliably and repetitively.

Since, the robot and humans were often working on the assembly of the same or neighboring car they often had to work in close proximity. We observed no visible change in the modus operandi of the human associates, when they were working along

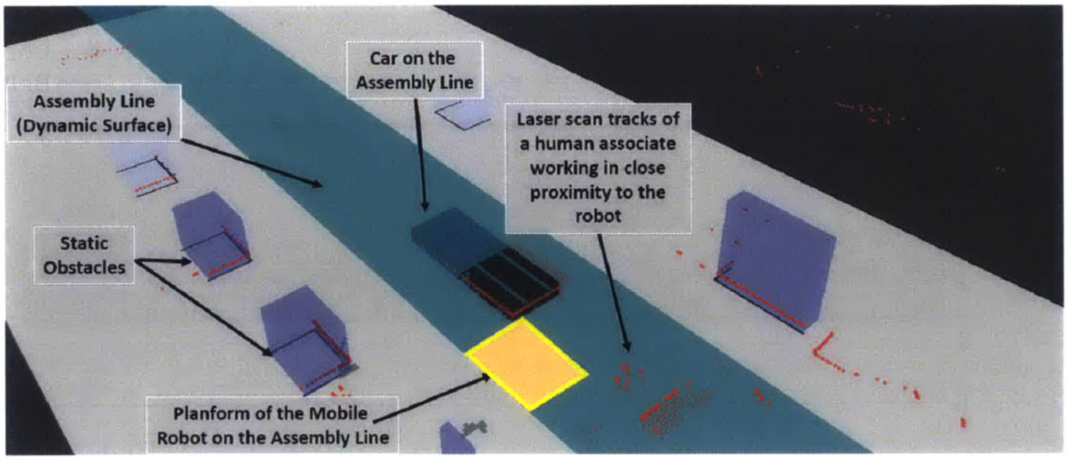

Figure 3-7: An rviz visualization of the mobile robot working in close proximity with human associates on the assembly of neighboring cars. The black box represents the car on the assembly line, and the red dots show the tracks of robot's laser scanner. The robot base is shown in yellow, and the conveyor belt of the assembly line is shown in dark green. We can observe the laser scan tracks of legs of a human associate working near the robot.

side the robot. As the mobile robot was monitored throughout the evaluation process, the observed interaction might not correspond to the interaction between human and an autonomous robot. However, this interaction does provide confidence regarding introducing robots right next to busy human associates.

3.3.7 Limitations and Future Directions

Along with testing the current system, the evaluation of the robot in an operational factory floor provided several directions for further improvements. In the current tests, the robot was continuously monitored by three engineers. This will not be true for the case when the robot is deployed in daily workflow of the factory. Thus, the system should be made more reliable, redundant and easy to maintain. On-board optic flow sensors are being used to detect transition between static and dynamic surface; one solution includes there use as an alternate measure of the surface speed, thereby increasing the redundancy in assembly line detection.

During certain runs the robot operation might be delayed due to the dynamic, uncertain events, thereby motivating the need for dynamic scheduling of factory tasks.

In addition, anticipatory path planning techniques and design interventions which enable fluent human-robot interaction may be used to reduce such delays as far as possible. Specifically, the time-critical nature of the task motivates the need to study and improve the interaction between busy, human associates and the autonomous, mobile robot.

Lastly, in the worst case scenario, a major challenge is to devise a fail-safe strategy to ensure human and equipment safety in the case when the robot stops functioning on the assembly line; though, this was not encountered in the current evaluation such a case cannot be neglected when the robot is working without any oversight during repeated daily operations.

3.4 Summary

This chapter presented the first published demonstration of a mobile robot capable of working directly on the automotive final assembly line in close proximity to busy, human associates. The designed system was demonstrated and evaluated in an operational factory floor, which included dynamic surfaces (conveyor belts), human associates, and other dynamic objects. Successful performance of the robot in these factory demonstrations, which was carried out over a period of five days, establishes confidence in the current system and paves the way for mobile robots that can collaborate on the dynamic surface of automotive final assembly lines. The design and evaluation also provided several directions for future work, which include improving system reliability and designing anticipatory path planning algorithms.

Chapter 4

Human-Robot Interaction in the Final Assembly

Robot autonomy is necessary but not sufficient for successful introduction of mobile robots on the automotive final assembly lines. The autonomous robot should also be able to work *fluently* along side the human associates in its surroundings. Small deficiencies in the human-robot interaction in a time-critical domain such as automotive final assembly can significantly degrade the efficiency of overall work-flow. This motivates the need for considering human factors while developing a robot for automotive final assemblies. In this chapter, we analyze human-robot interaction during the delivery phase of a repetitive fetch-and-deliver task. The development presented herein complements the algorithmic solutions developed in Chapters 2-3 towards the holistic design of an autonomous, collaborative robot for final automotive assembly.

4.1 Introduction

Human-robot interaction has been widely studied for multiple application domains, and is critical to successful introduction of robots among humans [31]. This holds true while designing interactive robots for the manufacturing domain that need to work in human-oriented environments. Akin to human-human collaboration, prior human-robot interaction studies have shown that a fluent collaboration requires awareness and anticipation of intent by both human and robotic agents [38]. To enable and facilitate this awareness and anticipation, it is important that both humans and robots communicate their status, location and intent, either explicitly through certain cues such as audio/visual signals or implicitly via their actions and motion paths. Our goal is to gain a better understanding of how such factors impact the efficiency and effectiveness of a robot assistant situated in an analogue assembly line environment, and suggest design interventions to improve the same.

As a first step towards improving the human-robot interaction in the final assembly line, we begin with analyzing the performance of the robotic assistant - both with and without design interventions - while collaborating with humans. Human-human interaction is often considered as a benchmark for human-robot interaction, thus we also compare the comparative performance of the robotic assistant to that of a human assistant. To do so, we design and conduct a human subject experiment in which we study these interactions in a controlled yet analogue assembly environment.

Leveraging its autonomous mobility developed in Chapters 2-3, the robotic assistant can perform myriad tasks on the factory floor. One of the primary tasks involving human-robot interaction will be that of delivering parts to human associates occupied with assembly tasks. Thus, the designed experiment focuses on a task involving delivery of parts by the assistant to busy, human associates.

Through the experiment, we assess objective and subjective measures of team fluency, assistant saliency, and investigate the requirement and effectiveness of explicit and implicit indicators designed to improve the human co-worker's awareness of the mobile robotic assistant. Specifically, we assess the effect of a flashing light on the robotic assistant and variations in the assistant's approach angle. Due to the presence of ambient noise in a typical factory floor, we consider a visual signal, i.e., the flashing light, as the explicit indicator instead of an audio signal. Prior work on human-aware navigation [64] suggests that humans prefer a robot path which is more visible, which motivates the use of approach angle as the implicit indicator. Lastly, we analyze whether a more *salient* assistant (utilizing the aforesaid indicators) results in a more *fluent* collaborator.

Outline of the Chapter

In this chapter, we begin with a brief overview of prior studies analyzing human-robot interaction in tasks similar to fetch-and-deliver. Next, we list the specific hypotheses we aim to evaluate. This is followed by a detailed description of the experiment design, protocol and evaluation measures. The statistical analysis of the experimental observations is presented next. Lastly, we discuss how the observed results inform the design of a more effective assistant, and also suggest interesting questions regarding robot saliency and its affect on collaboration, which warrant further investigation and analysis.

4.2 Related Work

Human-human collaboration during fetch-and-deliver tasks is seemingly natural, and does not require much cognitive effort for the human agents. For instance, while delivering or receiving an object from a fellow human being we are least concerned about the motion of our hands. However, much can be learnt about how to successfully carry out human-robot collaboration by studying human-human interactions [19, 40].

Prior studies on human-robot hand-overs provide useful insights for improving the motion of armed manipulators [23, 26, 67]. Experiments of give-and-take tasks between a human and robot [26], standing at a fixed locations, investigate the use of robot reaching gestures, vocalizations, and pauses as indicators for delivery. The study reports that communication using implicit natural actions is sufficient for give-and-take tasks, and the human collaborators do not require explicit instructions or expertise while collaborating with the robot. In other work, multiple studies on human-human and human-robot hand-overs result in design recommendations for robot motion, and formal descriptions of the physical and cognitive aspects of the hand-over process (see Fig. 4-1) [67]. Experiments in [23] investigate the use of contrasting motion to communicate the onset of a hand-over, and demonstrate statistically significant improvements in fluency of hand-overs by using temporal contrast during delivery. The study also reports a small but not significant increase in robot waiting time when the human participant is performing an attention task. Although a mobile robot is used in [23], the primary focus is on the motion of the armed manipulator, and modifications to motion that reduce the human waiting time.

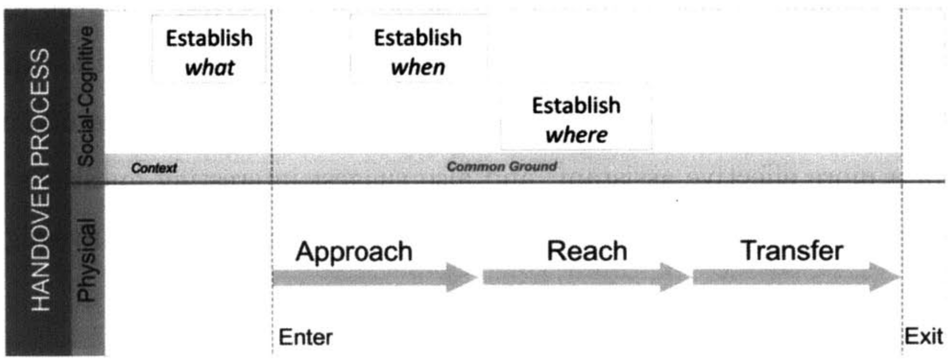

Figure 4-1: Formal description of the hand-over process [67]

Interaction between mobile robots and humans has been investigated for fetch-and-carry tasks. Studies evaluate robot approach direction and its impact on task efficiency and human comfort [43, 73]. Ideally results from these experiments would inform the design of motion planners that take into consideration human preferences. However, reported results are contradictory. Walters et al. [73] finds a left/right approach direction to be most favorable, while a rear approach direction to be the least favorable. In contrast, Koay et al.[43] indicates a frontal approach and delivery to be more favorable, especially for participants experienced with robots. These studies are carried out in settings without restrictions on approach directions, and with human participants that are primarily focused on the robot throughout its approach. Furthermore, no measurements concerning fluency of human-robot interaction have been made.

Though multiple studies evaluating robot's approach and fetch-and-deliver tasks have been conducted in the past, key differences exist in the domain of final assembly. A frontal approach is often not practical in a constrained factory environment. Oblique approach angles affect the human's visibility of the robot, and further study is needed into the effect of variations in robot approach angle from the rear direction. The robots are usually non-anthropomorphic, and the environment typically noisy. More importantly, human workers in a factory setting will usually be busy and not actively focused on the robot. Hence, there is a need to evaluate and improve human-robot interaction during the robot's approach towards an otherwise busy human co-worker.

4.3 Aim of the Experiment

Through human subject experimentation, we investigate interactions between an assistant and worker in an analogue assembly line environment, where the experiment participant takes the role of the worker. The task of the robotic or human assistant is to present the parts on a tray to the static human co-worker in a timely manner for continuation of the assembly task. Through the experiments we seek to evaluate the following hypotheses:

H1 *The interaction between the robotic assistant and worker during the delivery phase is less fluent than the interaction between a human assistant and worker. For this hypothesis, fluency is characterized by objective measures including workstation time and the assistant idle time.*

This hypothesis is founded in prior studies of hand-overs [67], which indicate hand-over quality degrades when working with a robot versus a human partner. We hypothesize that a similar effect exists, even when the assistant does not use manipulators and manipulation is the sole responsibility of the human worker. Idle times of agents (both human and robot) are indicative of fluency in a collaborative task, and have been used as objective measures of fluency in prior studies [38, 58, 51]. The design of our experiment ensures that the human worker is continuously occupied with tasks, and hence only idle times of the assistants (both robot and human) are evaluated. Interaction time, indicative of the total delivery time, is meant to quantify the time both the agents are interacting during the delivery phase of the task.

H2 *The worker subjectively assesses the interaction with the robotic assistant during* *the delivery phase to be less fluent than similar interaction with a human assistant. Subjective measures of fluency are evaluated using a series of Likert-scale questionnaires.*

Subjective measures of fluency are as important as their objective counterparts for evaluating human-robot collaboration. Hence, we evaluate the current hypothesis as a follow-on to **H1**. We have developed the questionnaire used in this experiment based on [37], which includes a survey of questions used to evaluate team fluency that produce values of Cronbach's alpha ≥ 0.7 (indicating measurement of similar latent variable).

H3 *Salient indicators for the robotic assistant improve the worker's awareness of the robot. Namely as indicators, we investigate the effect of variations in approach angle and the inclusion of a flashing light on the robot.*

Literature suggests awareness of the assistant and its intent improves task efficiency [38]. Factory settings are noisy, the workers' attention is occupied with assembly tasks, and the robot may not always be in the human worker's field of view. With this hypothesis we evaluate whether the specified indicators make the robot more salient. We measure the *look time*, the time of the participant's first head turn towards the assistant. An evaluation inspired by the Situational Awareness Global Assessment Technique (SAGAT) [27] is designed to measure the human worker's awareness of the robot in the task environment.

H4 *Salient indicators for the robotic assistant improve the objective and subjective measures of fluency for the robot.*

With this last hypothesis we evaluate whether the indicators, which may make the robot more salient, do indeed influence task fluency. Namely, we investigate the effect of indicators described in **H3**. The objective and subjective measures described in **H1** and **H2** are used to assess improvement in task fluency.

4.4 Experiment Methodology

The experiment is designed to simulate an analogue environment to the assembly line. In this setting the participant worker is standing at a workstation, stationary, and facing away from the assistant's approach path. The worker is provided an assembly task to occupy their attention. In the course of each trial two assistants, one robotic and one human, interact with the experiment participant. The assistants deliver parts enabling the continuity of the assembly task.

4.4.1 Materials and Setup

Figure 4-2: The Rob@Work mobile robotic platform augmented with additional sensors, a raised platform and tray. The modified setup is used as the robotic assistant for human subject experimentation.

We use a modified version of the Rob@Work robotic platform, shown in Fig. 4-2, as the mobile robotic assistant for this experiment. The basic platform is augmented with one Asus Xtion RGB-D device[1] used for person tracking, and one red, flashing and rotating light[2]. A raised platform and tray are mounted on the robot to represent the height of a future robotic arm. The robot, though capable of navigating autonomously, is operated manually by a human supervisor throughout the experiment.

[1]http://www.asus.com/Multimedia/Xtion_PRO_LIVE/

[2]http://goo.gl/yxreko

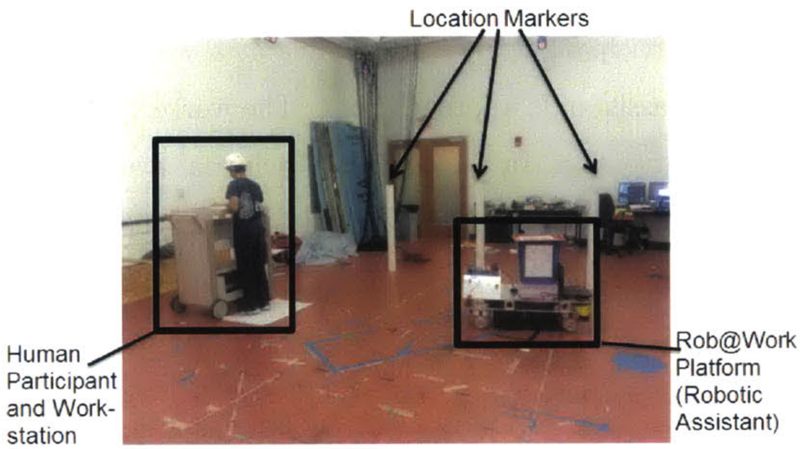

Figure 4-3: Experiment setup with human subject, workstation, Rob@Work platform, and white location markers.

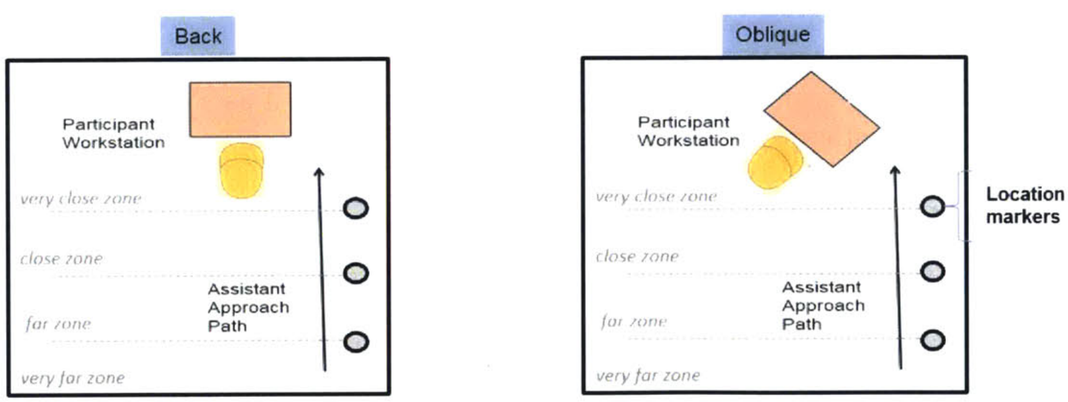

Figure 4-4: Schematic diagram of the experimental setup for the two approach angles (back and oblique).

Observations during the experiment are obtained using demographic survey, four in-experiment questionnaires, a post-experiment questionnaire, a Vicon motion capture system, video camera, and the on-board Xtion. The participants and the human assistant wear hard hats with Vicon markers. The robotic assistant also has Vicon markers mounted on its top to enable tracking via overhead Vicon cameras.

During each experiment, a recording of factory sounds[3] is played to simulate the factory environment. This serves to mask the noises made by the movement of the assistants to some degree, much like actual factory conditions would. Human participants are asked about the location of assistants in some questionnaires; thus, to eliminate the need to numerically estimate distances we place three large white poles equidistant from each other behind the participant. These poles are used as visual indicators to mark locations in the room (see Figs. 4-3-4-4), and divide the room into four names zones: *very close*, *close*, *far*, and *very far*, relative to the participant.

4.4.2 Procedure

During the experiment, each participant is instructed to stand at a specified location to work at a standing table as shown in Fig. 4-3. A dexterous, model-assembly task is presented to the participant. The task is chosen for its complexity and similarity to a factory assembly task, and involves constructing a Lego model (Fig. 4-5). Portions of the Lego parts and assembly instructions are delivered at specified intervals during the experiment by either the human or robotic assistant. Each delivery, but for the first, consists of three items: Lego parts, corresponding instruction set, and a questionnaire to assess level of awareness of the robot, and perception of safety, comfort, trust, and fluency. To keep the participants occupied prior to the first delivery, the first set of instructions (but not the Lego parts and questionnaire) are given during the briefing, and the first delivery contains only Lego parts and a questionnaire. The total experiment task time is ≈12-18 min, with a cumulative participant-assistant workstation time of $\approx 10 - 20$ seconds for each trial. This is similar to factory-like tasks, where delivery of parts constitutes only a portion of the overall assembly.

[3]`http://goo.gl/od11EQ`

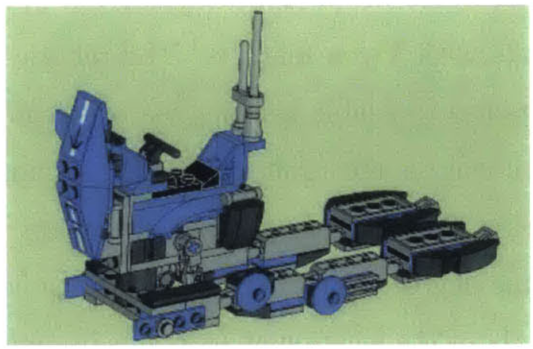

Figure 4-5: The Assembled Lego Model

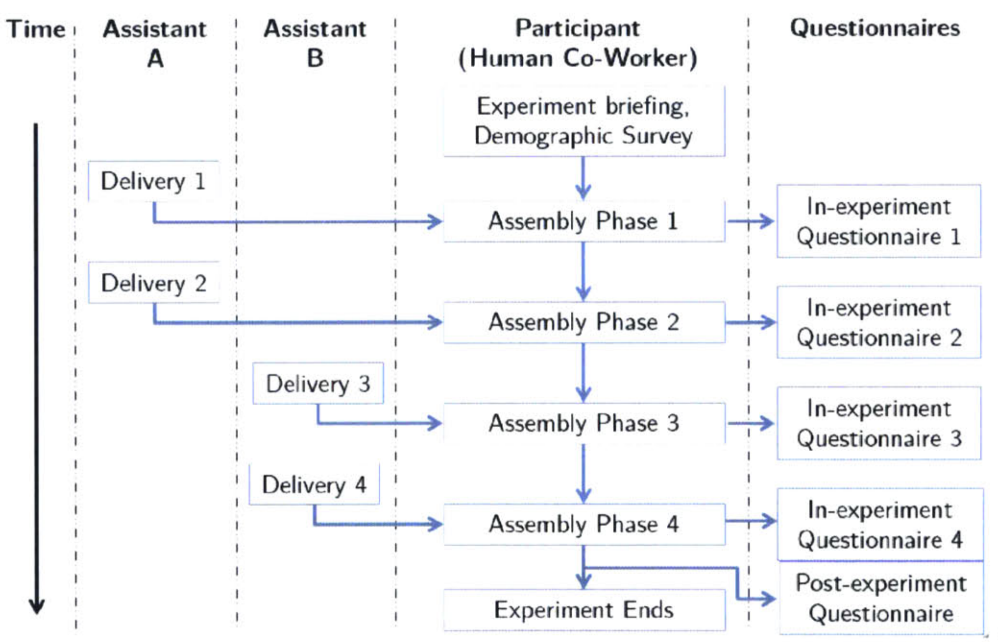

Figure 4-6: Timeline of the Experiment

An overview of the experiment's timeline is shown in Fig. 4-6. Each participant is aided by an assistant four times during the experiment, twice consecutively each by the human and robotic assistant (as shown in Fig. 4-7). This setup allows for each participant's response towards the robot to be compared to their own responses of working with a human assistant. Participants are randomly assigned to two groups one in which the first assistant is a human and the other in which the first assistant is a robot, to counterbalance differences that may arise due to the sequence of human and robot interactions. Further, trays are used by both human and robotic assistants during delivery, since the focus of the study is on the assistant's approach rather than manipulator motion.

Figure 4-7: Typical deliveries made by the Assistants (Top row: back approach, Bottom row: rear approach, Left column: robotic assistant, Right column: human assistant)

Participants are instructed to remove all the items from the tray once an assistant stops next to them. Further, they are told before the experiment that the assistants are involved in other time-sensitive tasks, to indicate the importance of the assistant's time. Using the delivered items, the participants build the part of the Lego model corresponding to the current instruction, and then answer the questionnaire. The assistant's approach towards the human co-worker, for the next delivery, starts as soon as the participant begins answering the previously delivered questionnaire.

When a delivery is not in progress, the human and robot assistants either move around behind the participants in the *far* or *very far* zones or stand still. These movements are scheduled for pre-determined times during the experiment, and are included to simulate the assistant's involvement in other factory tasks. At different, pre-determined times the appearance of the robot is altered by affixing different colored foam bumpers and printed signs; similarly, the human assistant alters his appearance by wearing/removing a safety vest, watch, and gloves. The color of the trays used by both human and robotic assistants, on which items are delivered, also varies between deliveries. The participants are not informed of these signs, indicators, and possibility of their modification a priori, since our goal is to evaluate the saliency of these features in a natural work-flow.

4.4.3 Design

The experiment is designed to evaluate the effect of three independent variables on the human-robot collaboration task:

IV1: Human or Robotic Assistant, varied within participants.

IV2: Robot's Flashing Light (Explicit Indicator) On or Off, varied between participants.

IV3: Assistant Approach Angle (Implicit Indicator) Back or Oblique, varied between participants.

The approach angle variable indicates the orientation of participant relative to that of the assistant's fixed approach path (Fig. 4-4). The back approach angle corresponds to the participant facing towards a wall and directly away from the assistant's path. The oblique approach indicates an orientation of approximately 45° away from the wall and towards the assistant's path. Additionally, participants are randomly assigned the first assistant as either the human or robotic assistant to counterbalance any learning effects and differences that may arise due to the sequence of human and robot interactions.

Table 4.1: Design of Experiment (N=24)

		Approach Angle		
		Back	Oblique	*Participants*
Light	On	H-H-R-R (3)	H-H-R-R (3)	12
		R-R-H-H (3)	R-R-H-H (3)	
	Off	H-H-R-R (3)	H-H-R-R (3)	12
		R-R-H-H (3)	R-R-H-H (3)	
Participants		12	12	**24**

Thus, the experiment is carried out as a mixed factorial design with four factors, one within- and three between-participants. Participants are randomly assigned across the 8 groups (2^k, where $k = 3$ is the number of between-participant factors) indicated in Table 4.1. The following sections describe the dependent measures observed for each participant.

4.4.4 Objective Evaluation

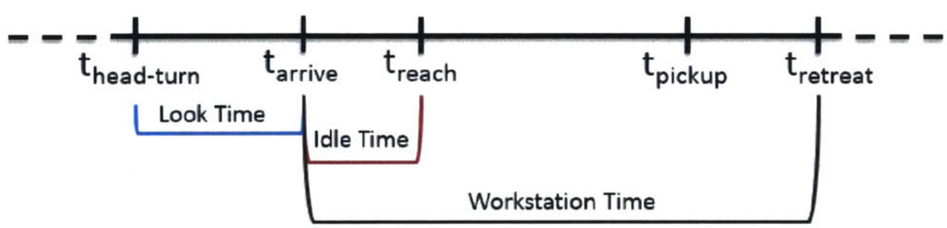

Figure 4-8: Key time epochs and derived measures based on the delivery phase of the experiment. Note that the order of these time instances and thereby sign of the time measures depends on the specific interaction.

Objective measures of team fluency are defined based on measures of time during the interaction. For the designed experiment, this interaction between the participant and the assistant (human/robot) primarily takes place during the delivery of objects. To extract the objective measures from the interaction, we identify key time instances during the interaction as shown in Fig. 4-8.

Interaction during a typical delivery begins when the participant first notices the approaching assistant. This measure is particularly difficult to identify, but can be approximated through either participant reported times or through certain codifica-

Table 4.2: Objective Measures: Definitions

Workstation Time	$t_{\text{retreat}} - t_{\text{arrive}}$ Time between assistant's stop on arrival and retreat, i.e., time required to complete the delivery.
Assistant Idle Time	$\max(0, t_{\text{reach}} - t_{\text{arrive}})$ Time between assistant's stop on arrival and the onset of the participant's reach towards the tray.
Look Time	$t_{\text{arrive}} - t_{\text{head-turn}}$ Time between participant's first head turn towards the assistant and the assistant's stop on arrival.

tion of the interaction. However, self-reporting interferes with participant's task and relies on participant's memory. Thus, we use the time when the participant first turns it heads towards the approaching assistant, denoted by $t_{\text{head-turn}}$, to approximate the time when the participant first notices the approaching assistant. Note that the participants are facing away from the assistant during the task, and need to turn their head only during the delivery. The assistant next arrives at the workstation, coded by the time instant t_{arrive}, and waits till the participant reaches for the objects on the tray, coded as t_{reach}. The time at which all the parts are picked from the tray is denoted as t_{pickup}. Once all the parts are retrieved by the participant the assistant retreats from the workstation at t_{retreat}. The order of the time instances shown in Fig. 4-8 is based on typical deliveries; however, certain interactions may involve a different order of $t_{\text{head-turn}}$, t_{arrive} and t_{reach}. For instance, the participant may turn his/her head towards the assistant after the assistant arrives at the workstation or the participant may reach for the objects on the tray before the assistant has stopped.

Based on the extracted time instances during interaction, we calculate the derived objective measures of team fluency; these include *workstation time* and *assistant idle time* and are inspired by previous studies of fluency in human-robot interaction [37]. Workstation time is defined as the difference between the assistant's stop on arrival and the beginning of the assistant's retreat, and measures the time of each delivery phase of the task. We also perform statistical tests using an alternate definition of workstation time given as $t_{\text{pickup}} - t_{\text{arrive}}$, to account for any confounding variations in robot's retreat that may occur due to teleoperation. No differences in statistical significance are observed for either of the definitions.

Assistant idle time is defined as the difference between the assistant's stop on arrival and the start of the participant's reach towards the tray. Idle time is a subset of the workstation time; we measure both since idle time specifically focuses on the time before which human actively interacts. In case the participant's reaching motion begins before the assistant arrives at the workstation, the idle time for that interaction is treated as zero.

Saliency of the assistant is quantified by the derived measure *look time*, the time between the participant's head turn towards the assistant and the assistant's arrival. Note that the *look time* may be negative or positive depending on whether the assistant is acknowledged with a head turn prior to or after its arrival. Table 4.2 defines these measures for the delivery task in the current experiment.

These measures are summed across the two deliveries, for each assistant, to obtain the cumulative measures reported in Section 5. The total task time, a usual measure of task fluency, is not used in this study since the assistants only deliver the parts and do not contribute in the actual assembly. This renders the total task time to be dependent on the participant's expertise in Lego assembly tasks and not on the interaction between the participant and the assistant.

These objective measures are independently coded by two raters from the video recordings of the experiment. The measured quantities are continuous and so Pearson's correlation coefficient is used to determine the inter-coder agreement. For all the derived quantities, the resulting correlation coefficient is ≥ 0.98 indicating very high inter-coder agreement. Degree of agreement in the coded data is further verified using the intra-class correlation coefficient [62] and the non-parametric measure Spearman's rho. Further, when determining statistical significance, data from both the coders is used independently to arrive at the final results, and any differences among coders are reported as errors.

To evaluate saliency of robot features, additional quantitative data is derived through in-experiment questionnaires. Participants are requested to fill out these in-experiment questionnaires between assembly steps. Questions are selected from the

Table 4.3: Objective Measures: Question Set

Awareness of Assistant
Common Questions
1. What color was the top surface of the tray on which the Lego pieces were delivered?
2. After delivering the parts, what was the robot/human assistant doing while you were working on the model? (i.e. which zone was he/it in, was he/it stationary or moving?)
Robotic Assistant
3. What colors were the bumpers on the robot?
4. What did the sign on the robot say the last time it delivered Legos?
5. Were the lights on top of the robot on during the last delivery?
Human Assistant
6. What color was the assistant's watch?
7. What color were the assistant's gloves?
8. Was the human assistant wearing a safety vest?

list shown in Table 4.3 pertaining to assistant's whereabouts and characteristics, and responses are evaluated as *correct* or *incorrect*. The questions pertain to the time immediately preceding the delivery of the questionnaire, and the features, of both the robotic and human assistant, examined via these questions are hidden from the participant while he or she answers the questions. This evaluation is inspired by the Situational Awareness Global Assessment Technique (SAGAT) [27], and focuses on Level 1 SA pertaining to "the perception of the elements [including status, attributes, and dynamics] in the environment within a volume of time and space."

4.4.5 Subjective Evaluation

Likert-scale statements and open-ended questions are used to collect subjective data regarding the experiment. Participants are administered a pre-experiment demographic survey and a post-experiment questionnaire. Further, subjective questions are also presented via the four in-experiment questionnaires. Several Likert-scale measures, derived from [37] and listed in Table 4.4, are used to evaluate participants' subjective response regarding comfort, safety, and perceived fluency. Some statements are repeated across questionnaires.

Table 4.4: Subjective Measures: Question Set

Comfort and Safety
1. I am comfortable with the time at which I first noticed the assistant.
2. I feel safe working in this environment.
3. I am comfortable with my level of awareness about the assistant's whereabouts.
4. I was stressed when the human/robotic assistant was assisting me.
5. I felt safe when working with the human/robot assistant.
6. I would have liked to notice the human/robot assistant coming earlier.
Fluency
7. The human/robotic assistant and I work well together.
8. Deliveries made by the human/robotic assistant were smooth.
9. I worked fluently together with the human/robot assistant.
Trust in Assistant
10. The human/robotic assistant's actions were consistent.
11. The human/robot assistant came when I expected him/it to.
12. The human/robot assistant's actions were predictable.
13. The human/robot assistant was dependable.
Additional Indicators
14. The human/robot assistant did his/its part successfully.
15. The human/robot assistant contributed to the success of the task.

4.5 Results

This section summarizes statistically significant results, trends, and other insights obtained from the experiments.

4.5.1 Participants

Thirty participants performed the experiment, out of which data from six participants could not be used in analysis due to incorrect task performance, non-completion of the experiment, or missing video data. The results presented in this section are for the 24 participants that successfully completed the experiment with complete data for analysis. These participants were randomly assigned amongst the three between-participant factors of the experiment, resulting in three replicates for each experimental setting (Table 4.1). The participants (13 men and 11 women) were recruited via email, and had a median age of 20 years (max = 31, min = 18). None of the participants indicated any form of colorblindness in the pre-experiment demographic survey. Prior to the experiment a pilot study was carried out with 5 participants to streamline the experimental procedure; data from the pilot experiment are not included in the reported analysis.

4.5.2 Task Time

Each participant took an average of 890 s ($\approx$ 15 minutes, SD=207.12 s) to complete the model-assembly task along with the in-experiment questionnaires.

4.5.3 Assistant Workstation and Idle times

A mixed factor, three-way Analysis of Variance (ANOVA) is carried out to compare the workstation times and assistant idle times across the four independent variables (1 within-, 2 between-participants) as detailed in Tables 4.5-4.6. Statistically significant differences ($p<0.05$, $N=24$) are observed in both these measures for use of human versus robotic assistant (the within-participant variable). No statistically significant

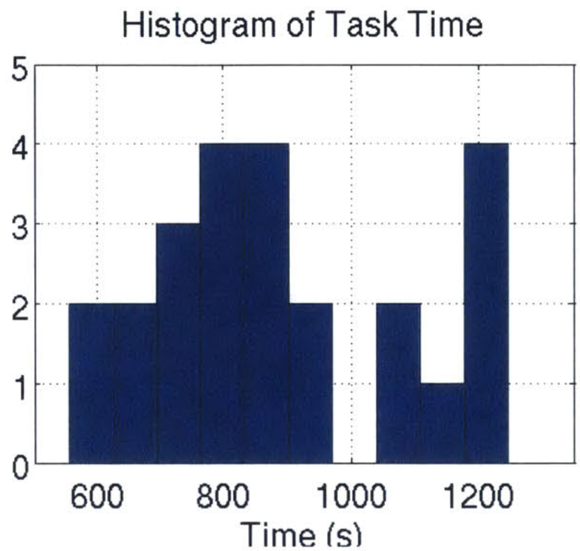

Figure 4-9: Histogram of Task Times

differences are found across the other factors, i.e., use of robot flashing light and approach angle. The alternate definition of workstation time yields similar statistical significance. To confirm that no learning effects are present, a four-way ANOVA was also performed which included type of first assistant as an independent variable. This randomization factor is found to be non-significant, while the other significance results are identical to the three-way ANOVA.

Table 4.5: Workstation Time: Mixed-factor Analysis of Variance

Source	SS	df	MS	F	p
Within-Subject Effects					
Assistant	84.748	1	84.748	6.699	**0.018**
A B	13.104	1	13.104	1.036	0.321
A C	2.134	1	2.134	0.169	0.686
A B C	10.773	1	10.773	0.852	0.367
Error	253.026	20	12.718		
Between-Subject Effects					
Angle (B)	20.254	1	20.254	0.729	0.403
Light (C)	13.632	1	13.632	0.491	0.492
B C	15.504	1	15.504	0.558	0.464
Error	555.571	20	27.779		

Figure 4-10 compares the mean for human and robotic assistants across all the participants for the objective metrics of fluency. Further analysis of the significant factors

Table 4.6: Workstation Time: Mixed-factor Analysis of Variance

Source	SS	df	MS	F	p
Within-Subject Effects					
Assistant	174.625	1	174.625	13.732	**0.001**
A B	37.395	1	37.395	2.941	0.102
A C	5.833	1	5.833	0.459	0.506
A B C	13.512	1	13.512	1.062	0.315
Error	254.339	20	12.717		
Between-Subject Effects					
Angle (B)	0.669	1	0.669	0.017	0.898
Light (C)	56.002	1	56.002	1.407	0.249
B C	0.002	1	0.002	0	0.995
Error	796.061	20	39.803		

confirms that the workstation times and idle times associated with the robotic assistant are statistically significantly higher than those for the human assistant ($p<0.05$, using two-tailed, paired t-tests to compare means with unknown, unequal covariance). The human assistant on average interacts with the participant for a cumulative time of 8.98 seconds (SE = 0.62 s) as compared to the robotic assistant's average of 12.8 seconds (SD = 1.31s). Similarly, the human assistant idles for a cumulative time of 1.45 second (SE = 0.43s), in contrast to the robotic assistant's idle time of 4.1 seconds (SE = 1.19s). These results support our first hypothesis **H1**, i.e., according to the objective measures of fluency the robotic assistant is a less fluent collaborator in comparison to the robotic assistant.

4.5.4 Subjective Measures of Fluency

Likert responses to statements about fluency (presented in Section 4.4.5) are analyzed across factors using non-parametric tests. No statistically significant differences are found in responses for the human versus robotic assistant (using the paired, Wilcoxon signed-rank test). Thus, the second hypothesis **H2** comparing the subjective measures of fluency is not supported by the experiment results.

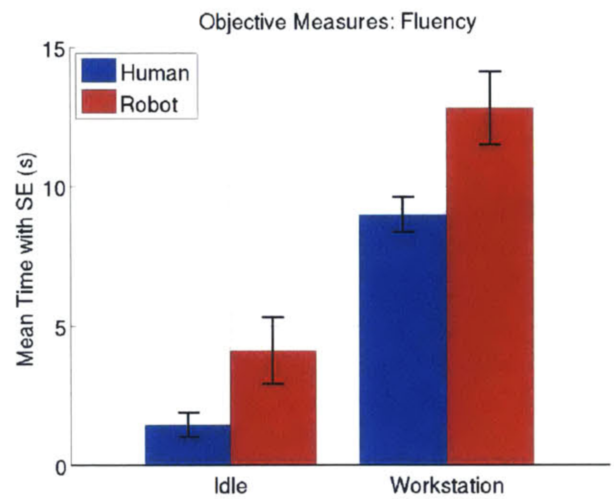

Figure 4-10: Objective Measures of Fluency: Differences across Assistant type

4.5.5 Saliency

Saliency of the assistant is evaluated using the derived measure *look time*. Responses to fourteen objective questions (listed in Table 4.3) pertaining to assistant's whereabouts and characteristics are also evaluated and frequencies of *correct* versus *incorrect* answers are compared. Overall no statistically significant effects are found for type of indicators, and the third hypothesis **H3** is unsupported. Type or use of indicators did not produce statistically significant differences in workstation and idle times, leaving the fourth hypothesis **H4** unsupported as well.

However interesting differences emerge based on type of assistant, human or robot (within-participant variable) as shown in Fig. 4-11. Analysis of look time is performed using a mixed factor ANOVA (see Table 4.7), and a statistically significant difference ($p<0.005$) is observed. The robotic assistant is noticed much earlier with average notice time of 8.82 seconds (SE=2.02s) prior to stop on arrival, as opposed to the human assistant's average notice time of 1.65 seconds (SD=0.71s). These results suggest that degradations in fluency are likely due to factors other than robot saliency, since on average the robot is acknowledged earlier than the human assistant.

Overall, there is only one statistically significant difference in the frequencies of correct responses between types of assistants. Color of the tray is noticed significantly

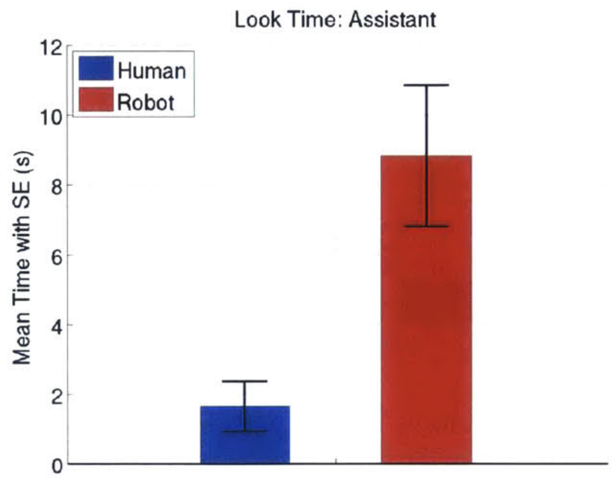

Figure 4-11: Objective Measure of Saliency: Differences across Assistant type

Table 4.7: Look Time: Mixed-factor Analysis of Variance

Source	SS	df	MS	F	p
Within-Subject Effects					
Assistant (A)	616.955	1	616.955	11.651	**0.003**
A B	0.201	1	0.201	0.004	0.951
A C	0.446	1	0.446	0.008	0.928
A B C	5.048	1	5.048	0.095	0.761
Error	1059.069	20	52.953		
Between-Subject Effects					
Angle (B)	48.736	1	48.736	0.728	0.404
Light (C)	31.698	1	31.698	0.474	0.499
B C	60.062	1	60.062	0.898	0.355
Error	1338.354	20	66.918		

better by the participants during the two deliveries each made by the robotic assistant (37%) than the human assistant (8%) ($p<0.05$, Fisher's exact test with 2×2 contingency table). However, participants demonstrated through responses on Question 2 (Table 4.3) that they were significantly more aware of their background environment after delivery by a human assistant ($p<0.001$). This suggests that the robot may have a transitory distracting effect that degrades situational awareness, even after the robot leaves the participant's side.

More visible features such as safety vest for human assistant and the state of the light for the robotic assistant were noticed by 79% and 67% of the participants, respectively. The participants were not informed about the existence or relevance of these signs in advance, and thus noticed them during the course of natural interaction with the assistant. Participants were equally unaware about less noticeable features such as the human assistant's glove color (13% correct responses) and watch (18%), and the robot's bumper color (4%) and printed signs (21%).

4.5.6 Factors Affecting Subjective Measures

The effect of flashing light and approach angle (between-participant factors) on Likert statement responses is evaluated using the two-sided unpaired Wilcoxon rank sum test. Participants agreed less strongly with the following statements when the robot's light was flashing, indicating a reduction in perception of safety and trust in the robot:

- I felt safe while working with the robot. ($p<0.05$, evaluated during post-experiment survey)

- The robot assistant's actions were consistent. ($p<0.01$ through in-experiment questionnaire, while not statistically significant during post-experiment survey)

Variation in approach angle resulted in statistically significant differences in responses as well. Participants agreed more strongly with the following statement in the 45° approach condition, indicating increased comfort with the robot when it approaches obliquely as opposed to from the rear (no such difference was found for the human assistant):

- I am comfortable with my level of awareness about the robot assistant's whereabouts. ($p<0.05$, assessed twice in the in-experiment questionnaires)

4.5.7 Open-ended comments

The participants are asked to provide open-ended comments about their experience with each assistant at the end of the experiment. Selected comments from these responses are included in Table 4.8. The open-ended responses suggest mixed reactions towards the robotic assistant. Interestingly, some comments reflect that participants felt rushed by the human assistant, and that the robot let them work at their own pace.

Table 4.8: Sample of Open-ended Comments

Human Assistant
"I liked that he said thank you!"
"Making eye contact & speaking was key to feeling more comfortable- I liked getting confirmation that the person wanted me to take stuff from the tray."
"Delivered parts when I'm still working, made me feel more stressed."
Robotic Assistant
"Smooth transition. Didn't get too close to me which I liked."
"Did a good job at a simple task."
"Having the robot moving around in the background was more distracting than the human in the background."
"With the robot, I think I made it wait till I'm done to get the stuff, I was less stressed."

4.6 Discussion

The results of the human subject experimentation support our hypothesis (**H1**) that the human-human collaboration is more fluent, as quantified by the workstation times and assistant idle times, than human-robot collaboration for the designed fetch-and-deliver task. Statistically significant results indicate that the robotic assistant spent on average 3.8 seconds (43%) more time than the human assistant interacting with the human for the same delivery tasks. Similarly, a statistically significant two-fold rise (2.65 seconds, 183%) is observed in the assistant idle time. A factory assistant will be making deliveries to multiple human co-workers. These differences in workstation and idle time, though small in magnitude, will substantially affect the productivity of the robotic assistant over the course of a two or three-shift work day, and should be alleviated using design interventions.

Surprisingly, the robotic assistant's approach is noticed on average much earlier (7.2 seconds, a statistically significant difference) by the participants as compared to that of the human assistant. Nonetheless, the robotic assistant idled more than the human assistant. This provides contrary evidence for the fourth hypothesis (**H4**), suggesting a salient agent does not necessarily produce an efficient collaborator.

These results suggest that degradations in fluency are likely due to factors other than robot saliency, since on average the robot is acknowledged earlier than the human assistant. While the participants reported the human and robotic assistant to be equally fluent, they appeared to be more comfortable with making the robot wait. This is supported by open-ended responses indicating different attitudes towards the human and robotic assistants. We posit that the human assistant's time is valued more than that of the robot, and personal objectives and comfort take a higher priority during collaboration with the robotic assistant. Further study is required to understand how to design the robot and its human interface so that it does not wait unnecessarily for human co-workers, which degrades the productivity of the robotic assistant and the overall assembly line workflow.

We confirm that an oblique approach angle is preferable since the participants report increased comfort about their awareness of the robot, although it did not improve the objective measures of fluency. The red flashing light is observed to be the most noticeable feature amongst those evaluated but did not improve objective measures of fluency. Further, the participants reported feeling less safe with the light flashing. This is possibly due to the color choice, and suggests that a red flashing light should not be used in nominal, safe operation of the robot.

4.6.1 Limitations

We designed our experiment to emulate a factory setting through careful choice of task and features such as noise, however limitations remain. The study was carried out in a large mostly empty room with student participants rather than factory workers. Further, the participants were working with the robotic assistant for the first time, and hence the effect of long-term experience working with the robot cannot be evaluated. In typical factory operations, the robot will be assisting multiple workers. Although participants were told that the assistant has additional responsibilities and tasks, it is possible human workers will behave differently when the robot's responsibilities clearly relate to other co-workers. Hence, there remains a need to study the interaction that include a longitudinal study with multiple co-workers, to observe how the interaction changes over time and with multiple people.

4.7 Summary

Successful introduction of mobile robots on the factory floor requires them to be capable of fluent human-robot interaction. To understand the factors that impact human-robot interaction, we conduct a human subject experiment to compare the performance of a mobile robotic assistant and human assistant. As fetch-and-deliver will be one of the primary tasks to be performed by the designed mobile robot, we use delivery of parts as the representative task to study human-robot interaction.

Results from the experiment indicate that interaction times and idle times are statistically significantly higher for the robotic assistant than the human assistant. However, the robotic assistant's approach is noticed on average much earlier by the participants as compared to that of the human assistant. These results suggest that degradations in fluency are likely due to factors other than robot saliency. Based on our observations, we conjecture that the human assistant's time is valued more than that of the robot, and personal objectives and comfort take a higher priority during collaboration with the robotic assistant.

We confirm that an oblique approach angle is preferable since the participants report increased comfort about their awareness of the robot. The robot's red flashing light did not improve objective measures of fluency, and the participants reported feeling less safe with the light flashing. This suggests that a red flashing light should not be used in nominal, safe operation of the robot.

The experiment provides initial guidelines for designing and improving human-robot interaction during delivery of parts in manufacturing domains. These insights will help make the autonomous robots not only a more efficient agent but a better collaborator while working along side humans on factory floors, including the automotive final assembly line.

Chapter 5

Conclusion

The primary objective of this thesis was to develop a mobile robotic system that can operate in automotive final assembly lines. We presented the desired robotic system that was designed by developing novel solutions, at both algorithmic and hardware level, and leveraging the prior art in mobile robotics. The robotic system was subsequently deployed in an operational automotive factory and its navigation performance was evaluated. Lastly, we analyzed the factors that impact the systems performance from a human-robot interaction perspective. This chapter briefly summarizes the key contributions of the thesis and discusses directions for future work.

5.1 Thesis Contributions

In summation, this section describes the key contributions of the thesis towards introducing mobile robots on the automotive final assembly lines.

Chapter 2: Trajectory Tracking on Dynamic Surfaces

Navigation on dynamic surfaces is a key prerequisite for introduction of mobile robots in automotive final assembly lines. To tackle this challenge we designed a modular position control algorithm that can track desired trajectories on both static and dynamic surfaces. The *modular* controller can be coupled with a variety of localization and path planning algorithms to enable robot navigation on dynamic surfaces.

Simulation using Gazebo robot simulator [44], off-the-shelf localization and planning algorithms, and assuming perfect knowledge of the surface speed demonstrate the efficacy of the designed control architecture in comparison to the nominal controller. In a simulated test scenario, where the robot has to traverse across a dynamic surface moving at 10 cm/s, the nominal control fails in tracking the desired trajectory while the designed control architecture results in a maximum deviation of 4 cm. In addition to tracking trajectories, by design the control algorithm avoids any adverse torques on the robot hardware due to the surface motion; an essential feature for maintainability of robot hardware in automotive factories.

In comparison to the nominal control architecture, the designed control architecture additionally requires as input the speed and location of the dynamic surface. We study various sensing alternatives, and design a sensing module capable of providing the required input signals to the robot in an automotive final assembly line. The sensing module comprises of four on-board optic flow sensors to detect the edge of assembly line, when the robot is transitioning to or from a dynamic surface, and an off-board contact-based sensor to measure the speed of the assembly line. A completely on-board solution using optic flow sensors is also discussed [71]; however, for robust performance in dynamic environments the final design consists of both on-board and off-board sensors.

The designed solution is implemented on the Rob@Work 3 robotic platform using off-the-shelf optic flow sensors and an off-board wheel encoder. The hardware implementation is tested in a controlled lab environment, wherein a customized treadmill is used to emulate the dynamic surface. As a first step, the position hold performance of the robotic base using the designed control architecture and sensing module is evaluated. The robot can successfully hold its position irrespective of whether it is completely on the moving surface or is straddling the same. A hardware test similar to the Gazebo simulation scenario is also performed, and the robot is observed to successfully traverse the moving surface. These software and hardware tests in static environments serve as a useful preliminary validation of the designed trajectory tracking solution, which is next evaluated in an operational automotive factory.

Chapter 3: Robot Evaluation on the Automotive Final Assembly Line

Equipped with the capability of navigating on dynamic surfaces, we deployed and evaluated the mobile robotic base in an operational, automotive final assembly line. To our knowledge this was the first instance of a mobile robot navigating dynamic surfaces of automotive factories. The mobile robot was tasked to perform proprietary assembly tasks on cars situated on an automotive final assembly line that was moving at an average speed of 8 cm/s. Two task scenarios were considered, which involved the robot to perform the assembly task on one or both sides of the moving car. Further, the task was to be performed along side humans in the dynamic environment and within the limited amount of time the car spent in robot's work area.

We report results from fifteen trial runs from the evaluation in the operational automotive factory. Using the designed trajectory tracking solution and custom localization and planning algorithms, the robot was able to successfully navigate in the operational, factory floor. The integrated, mobile robotic system took on average less than a minute to complete the time critical assembly task on one side of the moving car. Performance of the individual sensing and control sub-systems was also evaluated and found to be satisfactory. For the reported trial runs the on-board optic flow sensors were able to update the location of assembly line typically thrice for each of the trial run. The control sub-system was successfully able to track desired trajectories, and was independently validated using an additional, challenging position hold test while the robot was straddling the dynamic surface.

The successful evaluation of the mobile robot in the real environment establishes confidence towards introducing autonomous, mobile robots on the automotive final assembly line next to humans. During the evaluation, the mobile robot was monitored by three engineers; for autonomous operations the system should be made more robust, maintainable and user friendly. To be successful in time critical tasks, the robot additionally needs dynamic scheduling, anticipatory planning and interaction capabilities. Thus, factory evaluation suggests several directions for future work for successful introduction of robots among humans in dynamic, time-critical domains.

Chapter 4: Human-Robot Interaction: Collaborative Fluency and Robot Saliency

Factory evaluations motivate the need to improve the interaction capability of the mobile robot. In order to analyze the interaction between the mobile robot and busy humans, we designed and conducted a human subject experiment with 24 participants [72]. The experiment was conducted in an analogue assembly line environment, wherein the human associates (participants) had to perform an assembly task. The human associates worked with two assistants, namely, a human assistant and a mobile robotic assistant (Rob@Work 3). The task of the assistants was to deliver parts, which were required for the assembly task, to the human associate. By comparing the observed human-human interaction with the human-robot interaction, we aimed to study the factors that impact robot saliency and collaborative fluency. Both objective and subjective measures, motivated by prior research in human-robot collaboration, were used to measure robot saliency and collaborative fluency.

A statistically significant increase of 2.65 s (183%) was observed in the assistant idle time for human-robot interaction in comparison to that for human-human interaction ($p < 0.05$). This difference in objective measures of fluency suggests the need for design interventions to make the human-robot collaboration more fluent for time-critical assembly tasks. For the human and robot assistants, no statistically significant differences were observed in the subjective measures of fluency. Further, it was observed that the participants notice the robot (as measured by the measure: look time) earlier during the collaboration as compared to the human assistant. These results along with the open-ended comments from the participants suggest that degradations in fluency are likely due to factors other than robot saliency. Our experiment also evaluated the impact of robot's approach and a visual indicator (red flashing light). For these measures, no statistically significant differences were observed in objective measures. However, participants through the subjective measures reported increased comfort about their awareness of the robot for an oblique approach angle ($p < 0.05$). Next, we discuss few recommended directions for future work.

5.2 Recommended Future Work

The contributions of this thesis are important yet initial steps towards developing mobile robots that can work seamlessly with human associates for automotive final assembly. Automotive assembly present several additional challenges, both in research and implementation, that are beyond the scope of this thesis. For instance, issues of incorporating flexible robotic assembly in dynamic environments are discussed in [60]. Section 1.1.1 provides an overview of the implications of the dynamic environment of automotive final assembly for the mobile robot. Further, based on the factory evaluation of our system, Section 3.3.7 discusses the future steps for ease of integration of the mobile robot in the factory floor, especially where its autonomous operations will not be continuously monitored.

Here, we focus on two of these future directions, namely, efficient path planning in human-oriented environments and conveying the motion intent of robot. Future work in these two areas will make the robot navigation more efficient and intuitive to the robot's human collaborators, and complement the development of safe, autonomous navigation of robot in automotive final assembly lines presented in this thesis.

Path Planning among Humans

While developing the control and sensing solution in Chapter 2, we assumed that the desired robot path is available from a path planner or human teleoperator. Consequently, as a first step, in the factory evaluation of our robotic system we used an existing, reactive path planning approach [55].

However, this reactive approach creates plans based on the current sensing information and does not reason about the future locations of the dynamic objects and humans in the robot's surroundings. Further, the challenge of robot navigation while in proximity to humans includes aspects of physical human-robot interaction, wherein the robot must plan and execute trajectories toward goal locations that may interact with the trajectories of other agents in the environment. Prior work on the design of autonomous algorithms for robot navigation has shown that, in crowded environments, a robot benefits from anticipating the motion and cooperative behavior of humans [18, 47, 48, 70, 75]. Further, in absence of anticipation of the cooperative behavior, the robot stops frequently and exhibits the *freezing robot problem* [69].

Thus, we posit that in order to be effective, a robot navigating around humans must be capable of both predicting the motion of dynamic obstacles (including humans) and using this predictive information to plan safe and purposeful paths. This requires research in developing algorithms for (i) on-board detection of humans from mobile robots, (ii) real-time prediction of future states of humans, and (iii) anytime path planning using the anticipatory information of the environment.

Communicating Motion Intent

Research has shown that robot agents that are intent-expressive tend to collaborate better [38]. Further, navigation in human crowds is often a cooperative task in which humans implicitly collaborate [70]. For instance, we often give way to other pedestrians if we anticipate they are in hurry. As robot navigation amongst humans is a collaborative task and conveying intent improves collaboration, we believe, a robot which conveys its motion intent can navigate better among humans.

Prior work exists on legibility of robot motion, i.e. how to disambiguate the goals of the robot by purposefully modifying the robot motion [25]. This has led to the development of autonomous algorithms which aim to generate motion that conveys the intended goal and use the motion itself as the communication modality. Future work should explore design of systems that complement these algorithms and convey the planned path of robot using audio/visual communication modalities.

Bibliography

[1] Stefan Bartscher, Innovation Product Division, BMW Group, personal communication, February 2012.

[2] Adaptive Monte Carlo Localization - ROS Wiki. http://wiki.ros.org/amcl. [Online; last accessed 1-May-2015].

[3] Automated Guided Vehicle system at the automobile manufacturer Porsche. http://www.mlr.de/en/referenzen/automotive/porsche/. [Online; last accessed 1-May-2015].

[4] Baxter: manufacturing robot. http://www.rethinkrobotics.com/products/baxter/. [Online; last accessed 1-May-2015].

[5] BMW Group PressClub Global. https://www.press.bmwgroup.com/. [Online; last accessed 1-May-2015].

[6] InnoSenT miniature radars. http://www.innosent.de/en/industry/isys-systems. [Online; last accessed 1-May-2015].

[7] iRobot Roomba® vacuum cleaning robot. http://www.irobot.com/for-the-home/vacuum-cleaning/roomba.aspx. [Online; last accessed 1-May-2015].

[8] National Robotics Initiative (NRI). http://www.nsf.gov/funding/pgm_summ.jsp?pims_id=503641&org=CISE. [Online; last accessed 1-May-2015].

[9] Robots in manufacturing. http://www.britannica.com/EBchecked/topic/44912/automation/24853/Robots-in-manufacturing. [Online; last accessed 1-May-2015].

[10] Robots on BMW's Assembly Line. https://www.asme.org/engineering-topics/media/automotive/video-robots-bmws-assembly-line. [Online; last accessed 1-May-2015].

[11] Rob@Work 3. http://www.care-o-bot-research.org/robatwork-3. [Online; last accessed 1-May-2015].

[12] ROS Navigation for Care-O-Bot. http://wiki.ros.org/Robots/Care-O-bot/Tutorials/Navigation. [Online; last accessed 1-May-2015].

[13] Safety of human-robot systems in fixed workcell environments. `http://www.nist.gov/el/isd/ps/safhumrobsysfixedworkcellenvir.cfm`. [Online; last accessed 1-May-2015].

[14] Smart Robots Can Now Work Right Next to Auto Workers. `http://www.technologyreview.com/news/518661/smart-robots-can-now-work-right-next-to-auto-workers/`. [Online; last accessed 1-May-2015].

[15] VECNA hospital delivery robots. `http://www.vecna.com/solutions/hospital-delivery`. [Online; last accessed 1-May-2015].

[16] Prasad Akella, Michael Peshkin, ED Colgate, Wit Wannasuphoprasit, Nidamaluri Nagesh, Jim Wells, Steve Holland, Tom Pearson, and Brian Peacock. Cobots for the automobile assembly line. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 728–733. IEEE, 1999.

[17] I. Baldwin and P. Newman. Laser-only road-vehicle localization with dual 2d push-broom lidars and 3d priors. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, pages 2490–2497, 2012.

[18] Tirthankar Bandyopadhyay, Kok Sung Won, Emilio Frazzoli, David Hsu, Wee Sun Lee, and Daniela Rus. Intention-aware motion planning. In *Proceedings of the Workshop on Algorithmic Foundations of Robotics X (WAFR)*, pages 475–491. Springer, 2013.

[19] Patrizia Basili, Markus Huber, Thomas Brandt, Sandra Hirche, and Stefan Glasauer. Investigating human-human approach and hand-over. In *Human Centered Robot Systems*, pages 151–160. Springer, 2009.

[20] James C Boerkoel Jr and Julie A Shah. Planning for flexible human-robot co-navigation in dynamic manufacturing environments. In *International Conference on Human-Robot Interaction (HRI), HRI Pioneers*, 2013.

[21] Roger Bostelman and Will Shackleford. Improved performance of an automated guided vehicle by using a smart diagnostics tool. In *IEEE International Conference on Industrial Technology (ICIT)*, pages 1688–1693, 2010.

[22] Wolfram Burgard, Armin B Cremers, Dieter Fox, Dirk Hähnel, Gerhard Lakemeyer, Dirk Schulz, Walter Steiner, and Sebastian Thrun. Experiences with an interactive museum tour-guide robot. *Artificial Intelligence*, 114(1):3–55, 1999.

[23] Maya Cakmak, Siddhartha S Srinivasa, Min Kyung Lee, Sara Kiesler, and Jodi Forlizzi. Using spatial and temporal contrast for fluent robot-human hand-overs. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, pages 489–496, 2011.

[24] C.P. Connette, A. Pott, M. HaÌĹgele, and A. Verl. Control of an pseudo-omnidirectional, non-holonomic, mobile robot based on an ICM representation in spherical coordinates. In *Proceedings of the 47th IEEE Conference on Decision and Control, CDC 2008*, pages 4976–4983, Dec 2008.

[25] Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. Legibility and predictability of robot motion. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, pages 301–308, 2013.

[26] Aaron Edsinger and Charles C Kemp. Human-robot interaction for cooperative manipulation: Handing objects to one another. In *Proceedings of the International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1167–1172, 2007.

[27] Mica R Endsley. *Situation awareness analysis and measurement*, chapter Direct measurement of situation awareness: Validity and use of "SAGAT". Mahwah, NJ: Lawrence Erlbaum Associates, 2000.

[28] Dieter Fox, Wolfram Burgard, Frank Dellaert, and Sebastian Thrun. Monte carlo localization: Efficient position estimation for mobile robots. *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 1999:343–349, 1999.

[29] Dieter Fox, Wolfram Burgard, and Sebastian Thrun. The dynamic window approach to collision avoidance. *IEEE Robotics & Automation Magazine*, 4(1):23–33, 1997.

[30] Matthew Gombolay, Ronald Wilcox, and Julie A Shah. Fast scheduling of multi-robot teams with temporospatial constraints. In *Robotics: Science and Systems (RSS)*, 2013.

[31] Michael A Goodrich and Alan C Schultz. Human-robot interaction: a survey. *Foundations and trends in human-computer interaction*, 1(3):203–275, 2007.

[32] Birgit Graf, Ulrich Reiser, M Hagele, Kathrin Mauz, and Peter Klein. Robotic home assistant Care-O-bot® 3-product vision and innovation platform. In *IEEE Workshop on Advanced Robotics and its Social Impacts (ARSO)*, pages 139–144, 2009.

[33] Giorgio Grisetti, Cyrill Stachniss, and Wolfram Burgard. Improved techniques for grid mapping with rao-blackwellized particle filters. *IEEE Transactions on Robotics*, 23(1):34–46, 2007.

[34] Erico Guizzo. Three engineers, hundreds of robots, one warehouse. *IEEE Spectrum*, 45(7):26–34, 2008.

[35] Bradley Hamner, Seth Koterba, Jane Shi, Reid Simmons, and Sanjiv Singh. Mobile robotic dynamic tracking for assembly tasks. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2489–2495. IEEE, 2009.

[36] Bradley Hamner, Seth Koterba, Jane Shi, Reid Simmons, and Sanjiv Singh. An autonomous mobile manipulator for assembly tasks. *Autonomous Robots*, 28(1):131–149, 2010.

[37] Guy Hoffman. Evaluating fluency in human-robot collaboration. In *International Conference on Human-Robot Interaction (HRI), Workshop on Human Robot Collaboration*, 2013.

[38] Guy Hoffman and Cynthia Breazeal. Effects of anticipatory action on human-robot teamwork: Efficiency, fluency, and perception of team. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, pages 1–8, 2007.

[39] Dominik Honegger, Lorenz Meier, Petri Tanskanen, and Marc Pollefeys. An open source and open hardware embedded metric optical flow CMOS camera for indoor and outdoor applications. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2013.

[40] Markus Huber, Markus Rickert, Alois Knoll, Thomas Brandt, and Stefan Glasauer. Human-robot interaction in handing-over tasks. In *Proceedings of the 17th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 107–112. IEEE, 2008.

[41] Takayuki Kanda, Masahiro Shiomi, Zenta Miyashita, Hiroshi Ishiguro, and Norihiro Hagita. A communication robot in a shopping mall. *IEEE Transactions on Robotics*, 26(5):897–913, 2010.

[42] Oussama Khatib. Mobile manipulation: The robotic assistant. *Robotics and Autonomous Systems*, 26(2):175–183, 1999.

[43] Kheng Lee Koay, Emrah Akin Sisbot, Dag Sverre Syrdal, Mick L Walters, Kerstin Dautenhahn, and Rachid Alami. Exploratory study of a robot approaching a person in the context of handing over an object. In *Proceedings of the AAAI Spring Symposium: Multidisciplinary Collaboration for Socially Assistive Robotics*, pages 18–24, 2007.

[44] Nathan Koenig and Andrew Howard. Design and use paradigms for gazebo, an open-source multi-robot simulator. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2149–2154. IEEE, 2004.

[45] Stefan Kohlbrecher, Oskar Von Stryk, Johannes Meyer, and Uwe Klingauf. A flexible and scalable SLAM system with full 3d motion estimation. In *Safety, Security, and Rescue Robotics (SSRR), 2011 IEEE International Symposium on*, pages 155–160. IEEE, 2011.

[46] J Krüger, TK Lien, and A Verl. Cooperation of human and machines in assembly lines. *CIRP Annals-Manufacturing Tech.*, 58(2):628–646, 2009.

[47] Markus Kuderer, Henrik Kretzschmar, Christoph Sprunk, and Wolfram Burgard. Feature-based prediction of trajectories for socially compliant navigation. In *Robotics: Science and Systems (RSS)*, 2012.

[48] Frederic Large, Dizan Vasquez, Thierry Fraichard, and Christian Laugier. Avoiding cars and pedestrians using velocity obstacles and motion prediction. In *Intelligent Vehicles Symposium*, pages 375–379. IEEE, 2004.

[49] Steven M LaValle. *Planning Algorithms.* Cambridge university press, 2006.

[50] Sooyong Lee and Jae-Bok Song. Robust mobile robot localization using optical flow sensors and encoders. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 1039–1044. IEEE, 2004.

[51] Stefanos Nikolaidis and Julie Shah. Human-robot cross-training: computational formulation, modeling and evaluation of a human team training strategy. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, pages 33–40, 2013.

[52] Illah R Nourbakhsh, Clayton Kunz, and Thomas Willeke. The mobot museum robot installations: A five year experiment. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, volume 4, pages 3636–3641, 2003.

[53] Joelle Pineau, Michael Montemerlo, Martha Pollack, Nicholas Roy, and Sebastian Thrun. Towards robotic assistants in nursing homes: Challenges and results. *Robotics and Autonomous Systems*, 42(3):271–281, 2003.

[54] Morgan Quigley, Ken Conley, Brian Gerkey, Josh Faust, Tully Foote, Jeremy Leibs, Rob Wheeler, and Andrew Y Ng. ROS: an open-source Robot Operating System. In *International Conference on Robotics and Automation (ICRA), Workshop on Open Source Software*, 2009.

[55] Sean Quinlan and Oussama Khatib. Elastic bands: Connecting path planning and control. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 802–807. IEEE, 1993.

[56] G Reinhart and J Werner. Flexible automation for the assembly in motion. *CIRP Annals-Manufacturing Technology*, 56(1):25–28, 2007.

[57] Rolf Dieter Schraft, Christian Meyer, Christopher Parlitz, and Evert Helms. Powermate-a safe and intuitive robot assistant for handling and assembly tasks. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 4074–4079. IEEE, 2005.

[58] Julie Shah, James Wiken, Brian Williams, and Cynthia Breazeal. Improved human-robot team performance using chaski, a human-inspired plan execution system. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, pages 29–36, 2011.

[59] Jane Shi, Brad Hamner, Reid Simmons, and Sanjiv Singh. Mobile robotic assembly on a moving vehicle. In *ASME/ISCIE 2012 International Symposium on Flexible Automation*, pages 117–123. American Society of Mechanical Engineers, 2012.

[60] Jane Shi and Roland Menassa. Flexible robotic assembly in dynamic environments. In *Proceedings of the 10th Performance Metrics for Intelligent Systems Workshop*, pages 271–276. ACM, 2010.

[61] Jane Shi, Rick F Rourke, Dave Groll, and Peter W Tavora. Quantification of line tracking solutions for automotive applications. In *Performance Evaluation and Benchmarking of Intelligent Systems*, pages 311–338. Springer, 2009.

[62] Patrick E Shrout and Joseph L Fleiss. Intraclass correlations: uses in assessing rater reliability. *Psychological Bulletin*, 86(2):420, 1979.

[63] Roland Siegwart, Illah Reza Nourbakhsh, and Davide Scaramuzza. *Introduction to Autonomous Mobile Robots.* MIT Press, 2011.

[64] Emrah Akin Sisbot, Rachid Alami, Thierry Siméon, Kerstin Dautenhahn, M Walters, and S Woods. Navigation in the presence of humans. In *Proceedings of the 5th IEEE-RAS International Conference on Humanoid Robots*, pages 181–188. IEEE, 2005.

[65] Jean-Jacques E Slotine, Weiping Li, et al. *Applied Nonlinear Control*, volume 60. Prentice-Hall Englewood Cliffs, NJ, 1991.

[66] Siddhartha S Srinivasa, Dave Ferguson, Casey J Helfrich, Dmitry Berenson, Alvaro Collet, Rosen Diankov, Garratt Gallagher, Geoffrey Hollinger, James Kuffner, and Michael Vande Weghe. HERB: a home exploring robotic butler. *Autonomous Robots*, 28(1):5–20, 2010.

[67] Kyle Wayne Strabala, Min Kyung Lee, Anca Diana Dragan, Jodi Lee Forlizzi, Siddhartha Srinivasa, Maya Cakmak, and Vincenzo Micelli. Towards seamless human-robot handovers. *Journal of Human-Robot Interaction*, 2(1):112–132, 2013.

[68] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics.* MIT Press, 2005.

[69] Peter Trautman and Andreas Krause. Unfreezing the robot: Navigation in dense, interacting crowds. In *Proceedings of the International Conference on Intelligent Robots and Systems (IROS)*, pages 797–803. IEEE, 2010.

[70] Peter Trautman, Jeremy Ma, Richard M Murray, and Andreas Krause. Robot navigation in dense human crowds: the case for cooperation. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 2153–2160. IEEE, 2013.

[71] Vaibhav V Unhelkar, Jorge Perez, James C Boerkoel Jr, Johannes Bix, Stefan Bartscher, and Julie A Shah. Towards control and sensing for an autonomous mobile robotic assistant navigating assembly lines. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, 2014.

[72] Vaibhav V Unhelkar, Ho Chit Siu, and Julie A Shah. Comparative performance of human and mobile robotic assistants in collaborative fetch-and-deliver tasks. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, 2014.

[73] Michael L Walters, Kerstin Dautenhahn, Sarah N Woods, and Kheng Lee Koay. Robotic etiquette: results from user studies involving a fetch and carry task. In *Proceedings of the International Conference on Human-Robot Interaction (HRI)*, pages 317–324, 2007.

[74] Peter R Wurman, Raffaello D'Andrea, and Mick Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9, 2008.

[75] Brian D Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, James A Bagnell, Martial Hebert, Anind K Dey, and Siddhartha Srinivasa. Planning-based prediction for pedestrians. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 3931–3936. IEEE, 2009.