Copyright

by

Rolando Fernandez Jr.

2018

The Thesis committee for Rolando Fernandez Jr. certifies that this is the approved version of the following thesis:

Light-Based Nonverbal Signaling with Passive Demonstrations for Mobile Service Robots

APPROVED BY

SUPERVISING COMMITTEE:

____

Peter Stone, Supervisor

____

Justin W. Hart

Light-Based Nonverbal Signaling with Passive Demonstrations for Mobile Service Robots

by

Rolando Fernandez Jr.

Thesis

Presented to the Faculty of the Graduate School
of the University of Texas at Austin
in Partial Fulfillment
of the Requirements
for the Degree of

Master of Science in Computer Science

The University of Texas at Austin
May 2018

Acknowledgments

First, I would like to thank Piyush Khandelwal, the Building-Wide Intelligence (BWI) project members, and everyone else in the Learning Agents Group (LARG) who contributed to making the BWIBots what they are today, and without whom this thesis would have surely not have been possible. I also thank them for their input on my research as well as the presentation of my work.

I would like to thank Nathan John and Sean Kirmani for their assistance with my experiments and user studies.

I am grateful to all the participants who volunteered their time to participate in my user studies.

Thank you, Patricio Lankenau, for encouraging the use of programmable light strips on the BWIBots.

Finally, I express my sincere gratitude to my advisor Prof. Peter Stone and Dr. Justin W. Hart for their continuous feedback and their helpful insight throughout my research.

This work has taken place in the Learning Agents Research Group (LARG) at the Artificial Intelligence Laboratory, The University of Texas at Austin. LARG research is supported in part by grants from the National Science Foundation (IIS-1637736, IIS-1651089, IIS-1724157), Intel, Raytheon, and Lockheed Martin.

Light-Based Nonverbal Signaling with Passive Demonstrations for Mobile Service Robots

by

Rolando Fernandez Jr., MSCompSci
The University of Texas at Austin, 2018
Supervisor: Peter Stone

With emerging applications in robotics that have the potential to bring them into our daily lives, it is expected for them to not only operate in close proximity to humans but also interact with them as well. When operating in crowded, human-populated environments there are many communication challenges faced by robots due to variable levels of interactions (e.g. asking for help, giving information, or navigating near humans). A crucial factor for success in these interactions is a robot's ability to express information about their intent, actions, and knowledge to co-located humans. Many of the robot platforms developed for service roles have non-anthropomorphic form factors in order to simplify and tailor them to their jobs. Due to a lack of anthropomorphic features, these types of robots primarily communicate using an on-screen display and/or spoken language. To overcome the limitation of not communicating as people do, we explore the viability of nonverbal light-based signals as a communication modality for mobile service robots. These types of signals have many benefits over existing modalities which they can either complement or replace when appropriate, such as having long-range visibility and persisting over time. We present a novel light-based signal control architecture implemented as a custom Robot Operating System (ROS) software package generalized to allow for various signal implementations. We implement our framework on a BWIBot, an autonomous mobile service robot created as part of the Building-Wide Intelligence Project, and evaluate its validity through a real-world user study on the scenario where a robot and human are traversing a shared corridor from opposite ends, and the potential conflict created when their paths meet. Our results demonstrate that exposing users to the robot's use of an animated light signal only once prior to when it is information critical for the user is sufficient to disambiguate its meaning, and thus greatly enhances its utility in-situ, with no direct instruction or training to the user. These findings suggest a paradigm of passive demonstration of light-based signals in future applications.

Table of Contents

List of Tables

List of Figures

1

Introduction

One vision for mobile service robots is that they will be able to assist humans in a variety of scenarios and environments by performing a complex assortment of tasks to support day-to-day work, domestic, and care activities. Service robots are expected to provide assistance in domains including the home (6, 13, 32, 36), space exploration (31), schools, and workplaces (6, 20, 29, 32, 36, 41). These robots are often constructed in form factors and appearances that are specialized to the tasks which they perform. Many have a functional, non-humanoid appearance with few to no anthropomorphic features. Gracefully interacting with the people who use these devices will involve studying interaction paradigms that differ from the more well-studied areas of human-robot interaction that are based on humanoid or android machines (5). Communicating a variety of unique and recognizable signals in situations where humans cannot be trained or where the robot is not directly interacting with any specific person remains challenging (3, 4, 12). Common communication modalities available to service robots, such as spoken language and on-screen displays, are limited by proximal constraints and can be ineffective when the human is too far to see the display or hear the robot's voice (3).

The Building-Wide Intelligence (BWI) project at The University of Texas at Austin seeks to develop mobile service robots, called BWIBots (Figure 1.1), that can provide assistance to the occupants and visitors of the Computer Science Department (29). This complex environment presents many obstacles to a robot freely navigating its hallways, which can be overcome through coordination or collaboration with its human occupants (33, 41). One such challenge our BWIBots face while autonomously navigating the corridors of the department throughout the workday is the frequent undesirable occurrence of a robot and human blocking each other's path. When conflicts arise in our lab area, we often find them humorous, as these robots are research prototypes. If these robots were deployed in mission-critical applications such as hospitals or even less critical applications such as delivering room service in hotels, the problem of creating such a blocked passage would escalate from being a humorous chance event to being a critical design flaw. To develop solutions to this problem, we have constructed a test hallway (Figure 1.2) in a large laboratory space where we study various navigational scenarios in which humans and robots must negotiate shared space. The work presented in this paper considers the scenario of a robot and a human navigating a corridor from opposite ends in opposing directions, and the potential conflict created when their paths meet. Throughout this work, we will focus our analysis on a custom-built autonomous mobile service robot, the BWIBot, which performs different types of tasks in a office like environment. The experiments presented in this work are designed such that the findings are applicable to a variety of robot platforms, tasks, and applications.

This work explores two concepts in non-anthropomorphic Human-Robot Interaction (HRI). The light-based signal system, as designed, uses a pair of light emitting diode (LED) light strips mounted to the robot's chassis to indicate the robot's intended motion trajectory in a fashion similar to the turn signals on a car. The robot's navigation algorithm treats the corridor as being divided into three traffic lanes through which it may navigate. When "changing lanes," the robot

Figure 1.1: The BWIBot family of robots. Left: V2 with Arm, Center: V3 (the robot used in this study), Right: V2.

Figure 1.2: Constructed hallway environment with robot and participant in the early stage of hallway traversal.

signals this "lane-changing" behavior by blinking the LED light strip on the side of its chassis matching that of the direction of the lane that it intends to shift into. The robot can change lanes in order to avoid navigating into a path which conflicts with that of the person, however, the hallway is narrow enough that the person must also change into the opposing lane, in order to provide sufficient space for conflict-free passage. The first concept we explore in this study is the naturalness and efficacy of this signal in indicating the robot's intention.

The second concept is that of passively demonstrating the signal by having the robot use it in a context where the person can witness its usage, but prior to an interaction that necessitates understanding the signal's intent. This paradigm of establishing the meaning of the signal in advance of its necessity enables the person to be passively introduced to the signal's intention without requiring explicit, up-front training. The idea to introduce a passive demonstration to our interaction comes from a pilot study in which 13 participants (9 male, 4 female) reacted ambiguously to the LED signal; attempting to change lanes, but with some interpreting it as an instruction regarding which lane to shift into and others interpreting it as indicating the robot's intended path. In the study presented in this paper, the robot and the human traverse a 17.5 meter long corridor from opposite directions, passing each other along the way. To demonstrate its lane-changing behavior, the robot performs a lane-change in front of the user at the start of its path during this interaction.

These concepts are evaluated in a 2×2 user study, the conditions of which are described in Table 1.1. On the first axis, the robot either does or does not employ the LEDs to signal its lane-change behavior. On the second axis, the robot either does or does not perform a passive demonstration to the user prior to com-

Study Conditions			
		Passive Demonstration	
		No Demonstration	**Demonstration**
LED Signal	**No LED**	No Demonstration, No LED	Demonstration, No LED
	LED	No Demonstration, LED	Demonstration, LED

Table 1.1: Study Conditions

ing close enough to the user to force the robot to change lanes in order to avoid a conflict. The results of this study demonstrate only a modest improvement in reducing the number of navigational conflicts in the case of using the LED signal in the lane-changing behavior, though it becomes a very large improvement when the passive demonstration is introduced. Our findings suggests that user understanding of nonverbal robotic signals can be significantly enhanced simply by allowing users to witness them being used in their context.

1.1 Motivation

The research presented in this thesis is inspired by a two practical goals.

Goals:

1. Implement a system that enables a mobile service robot to use light-based nonverbal signals as a communication modality.

2. Evaluate whether light-based signals are a viable modality for conveying information to co-located humans.

1.2 Contributions

The primary contributions from this work are:

1. A *light-based signal control architecture* implemented in ROS, which controls addressable light strip arrays.

2. A *real-world user study* investigating the viability of light-based signals as communication modality for a mobile service robot.

3. The proposal of a *paradigm of passive demonstration of light-based signals* for future applications on robots.

2

Related Work

This thesis will be focusing on the use of light-based signals as a communication modality for mobile service robots. In this chapter, we hence first introduce a short overview of the general uses of lights in different applications and how they have previously been used as a communication modality. We then focus our review on the use of lights on robots for communication which will be the subject of Chapter 4.

2.1 Uses of Lights

Light signals have long been used to communicate information across long distances and in limited visibility environments. Though light-based communication is simple and only able to communicate very little information content when compared to other visual modalities, such as on-screen displays, it can be easily perceived and vary in levels of intensity (9, 11, 26). This allows for light-based signals to be used as a communication modality across different levels of notification and criticality (5).

Examples of light communication can be seen in aviation and maritime navigation (21). Historically, the light signals used in these applications were usually very complex and required prior training to be understood, with one well-known example being Morse code (27). Lights produced though bioluminescence also play part in the communication of some animals species (e.g., jellyfish, phytoplankton, fireflies, and pyrophorus) for purposes such as passive defense, baiting prey, or finding a mate (25).

A commonly-recognizable use of light signals in society is in the domain of automotive navigation. Light signals are used to alert drivers to the actions of others on the road in the form of brake lights and turn signals, in traffic lights to control the flow of traffic, and to alert drivers to emergencies. Additionally, they are commonly used to convey emergency information in buildings, such as with fire alarms.

2.2 Lights as a Communication Modality

Dynamic visual cues created with lighting displays have been shown to have the ability to convey complex properties, such as animacy and intent, even in simple animations (7). Further work with lighting displays has revealed that light can even be used elicit complex social responses to include emotion (8). Since mobile service robots tend to have more functional appearances, configurable multi-color light arrays are a natural choice of signaling mechanism.

Addressable RGB LED strips on which we focus on in this thesis, allow for fine control of multiple LEDs with the ability to change the color and intensity of each individually. Each LED on the strip is a combination of an LED and a micro-controller that allows for the control of color intensity as a function of time. Through the propagation of serial communication packets across an LED strip, variety of dynamic light animations can be created using a one-wire interface.[1]

[1] https://github.com/pololu/pololu-led-strip-arduino

2.3 Light-Based Robot Communication

Despite their abundance in society, research on the use of lights on robots for communication is still in its early stages by comparison to research involving humanoid robots and human-like behaviors (5). Some current applications concentrate on expressing information such as subtle expressions (19) or simulated emotions (18) rather than directly communicating instructions or important state information. Others serve very basic functions which descend from consumer electronics, such as reporting the robot's battery status and or powered on state (1, 5).

Other research has explored how lights can be used to communicate a robot's internal state and functional intent (5). Lights have been used to represent which actor is currently attempting to speak in a Human-Robot Interaction (HRI) dialogue task (9). Animated lights have also been used to give aerial drone robots the ability to communicate navigational intent (38). Autonomous motor vehicle applications for animated lights have also been explored to communicate the direction of turning and when it is safe to cross in front of the vehicle (37).

Most relevant to our work, Baraka et al. (1, 2) developed a framework for using animated lights for navigation tasks such as going from one location to another, as well as for human interaction tasks such as asking for human help. For navigation tasks, they used LEDs to indicate information such as percent done for an escort task, being blocked by an obstacle, or whether the robot is turning left or right. However, a limitation of their experiments is that participants were asked if they could interpret the meaning of the lights in a video of their robot performing various tasks. In contrast, we examine how people interpret LED signals on robots in the context of a natural and common human-robot interaction scenario, that of passing a robot going in the opposite direction in a hallway.

3

Design and Implementation of Light-Based Signaling System

In this chapter, we provide a ROS implementation of a system which allows for light-based signals to be used as a communication modality on mobile service robots. This system will facilitate the process of communicating with a LED-based signal mechanism to convey information about a robot to co-located humans. We begin by describing the BWIBot platform which we use for this thesis. We then present the design of the LED-based signal mechanism and the software that facilitates communication with this mechanism. Our implementation can be easily extended to other LED arrays and animations besides those discussed.

3.1 Building-Wide Intelligence Project

Our light-based signal control architecture is implemented on the BWIBots of the BWI project, which has the overarching goal of developing fully-autonomous mobile service robots that are able to exist as permanent entities in the Computer Science department of The University of Texas at Austin (29). The BWIBot robot platform is a custom robot platform that has been in development for over five years. Currently, our lab has four robots in operation and one in development. Three of the current robots are based on our second generation platform, using older hardware and software (Figure 3.1a). One of our second generation robots is also equipped with a Kinova Mico Arm[1] (Figure 3.1b). The fourth robot is our

[1]The Kinova arm depicted has 6 degrees of freedom, a reach of 700 mm, and a mid-range payload capacity of 2.1 kg

(a) Version 2 (b) Version 2 with Arm (c) Version 3

Figure 3.1: The three current BWI robotic platforms

only robot based on the third generation platform design (Figure 3.1c). The fifth and latest robot is a the development platform for our fourth generation platform, it is quite similar to the third generation but has a updated base and no center LiDAR.

Leveraging the BWIBot platform our lab has published extensive work in the field of Artificial Intelligence and Robotics. One of the our key areas of research is planning, particularly hierarchical motion planning for navigating to different locations (16, 42) and multirobot symbolic planning that allows for a team of our robots to achieve their individual goals while minimizing the overall cost of the plan (34). Using the multiple BWIBots our lab has researched multi-robot guidance that would allow our robots to work in concert to guide humans to their desired destinations (15, 17, 28). We have also developed a novel framework for evaluating the capabilities of intelligent mobile robots called "Robot Scavenger Hunt"

(44). With the addition of the Kinova arm we conducted research grounded language learning, such as by playing a game of "I Spy" (40), through human-robot dialog (39), and active learning (14). The Kinova arm also allowed for our lab to develop a novel framework for object exploration and ordering using haptic exploratory behaviors (35). "Virtour" a telepresence system allowing remote tours of our lab in the Computer Science department through the BWIBots has also been implemented (22).

Our light-based signal control architecture is able to support all three generations, only needing to have the LED segment parameters adjusted for each specific platform due to variation in height. However, at this time we have only installed LED strips on our third generation platform, thus we will only describe the hardware for this platform generation in detail.

3.1.1 Hardware

All of the robots in the BWI project are built upon the Segway Robotics Mobility Platform (RMP) created by Stanley Innovation.[2] Our third generation robot is built upon a custom prototype version of the newer RMP 210 model,[3] which comes with two integrated lithium-ion batteries, can navigate at speeds up to 8 m/s, and carry up to 45 kg of payload. The frame of the robot was designed in our lab and is capable of supporting a large array of sensors and devices. For navigation, localization, and obstacle avoidance, we utilize a Velodyne VLP-16 Puck LIDAR.[4] For video data, we use an ASUS Xtion PRO Live RGB-D Camera,[5] which

[2]http://stanleyinnovation.com/products-services/robotics/robotic-mobility-platforms/

[3]http://stanleyinnovation.com/products-services/robotics/robotic-mobility-platforms/passive-stability/

[4]http://velodynelidar.com/vlp-16.html

[5]https://www.asus.com/us/3D-Sensor/Xtion_PRO_LIVE/

also provides point cloud data (3D voxel maps). The third generation robots also have an Hokuyo URG-04LX laser range finder[6] to compensate for the Velodyne's near-field blind spots. The core computer system is a custom-built compact fanless computer[7] which runs Ubuntu 16.04 64-bit. The computer and all other devices on the robot are powered by the RMP 210's dual lithium batteries, eliminating the need for an external lithium battery (which is utilized in our second generation robots). The battery life of the third generation robot is approximately 6 hours when actively running and 10 when stationary.

3.1.2 Software

The software architecture for our BWIBot platform is implemented ROS (30). ROS provides the infrastructure for running a distributed node robot system, as well as the messaging framework required to communicate between the different nodes. Furthermore, as open source software ROS allows for robot users to concentrate on their core research interests by providing access to many packages and tools, such as device drivers, navigational systems, and planning systems.

The navigational system in our robot platform utilizes a hierarchical task-planning architecture for navigation planning (43). When a navigational action is requested it begins at the logical planner, which uses the Answer Set Programming (ASP) language to describe the operational environment (24) – such as which rooms are on which floors, and which doors belong to each room – and then processes the ASP logic using the ASP solver Clingo to generate all the possible plans

[6] https://www.hokuyo-aut.jp/02sensor/07scanner/urg_04lx.html

[7] 3.9Ghz i7 processor, 16GB DDR3 RAM, Gigabyte GA-Q87TN motherboard, HDPLEX H1.S heatsink case, Acer FT200 monitor, and a Logitech K400 Plus Wireless Touch Keyboard with Built-in Trackpad

(10) that satisfy the navigational request. It then randomly chooses one of the shortest plans and proceeds to the logical navigator which uses current and previous sensor readings (in the form of occupancy maps) and it's current knowledge of the environment to create a navigational path plan. Lastly, the local planner utilizes the immediate sensor readings while communicating movement commands to the segway base contoller while actively avoiding any obstacles that are encountered.

A large portion of the software for BWIBots comes from open source ROS packages (e.g. for navigation we use the move base ROS package). Furthermore, all of the code that is written by our BWI lab is open source as well and available to the public on github[8]. Some of our software repositories are released as ROS packages to the community and are available for install as binaries[9].

3.2 LED Control Server

We now provide a detailed description of the LED-based signal control system implemented for this thesis. We cover both the hardware of the physical LED-based signal mechanism and the software that facilitates communication with this mechanism.

3.2.1 Hardware

To be able to use LED animations on our BWIBots, a mechanism which could be mounted to the current frame of the robot platform with minor modifications was designed and prototyped. The mechanism consists of an array of 120

[8] https://github.com/utexas-bwi

[9] http://wiki.ros.org/bwi, http://wiki.ros.org/bwi_common, and http://wiki.ros.org/segbot

individually-controllable, multi-color LEDs constructed from (2) Pololu Addressable RGB 60-LED Strip, 5V, 2m (WS2812B)[10] (Figure 3.2c) in combination to provide coverage of both the front and rear of the robot. The LED strip enables the manipulation of three variables: LED color (RGB [Red, Green, Blue]), intensity [0,255], and position [0,119]. To control the LED array, an Adafruit Metro Mini 328 - 5V 16MHz[11] (Figure 3.2a) micro-controller with the capability of communicating over serial using a Universal Serial Bus (USB) connection was added to the robot. Furthermore, to power and interface the LED strips with the micro-controller a compact printed circuit board (PCB), which allows for the LED strips power and communication connections to be centrally connected, was created. The PCB, UTexas BWI LED Power Board[12] (Figure 3.2b), was designed using the Fritzing[13] software and printed through the company OSH Park[14].

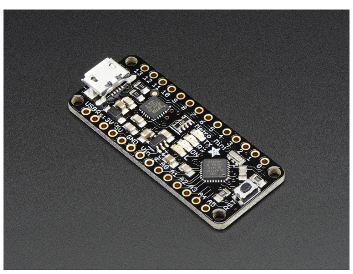

(a) Adafuit Metro Mini
Source: Adafruit.com

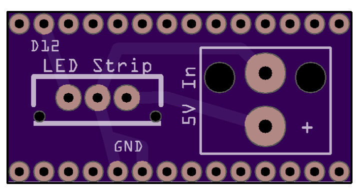

(b) LED Power Board PCB
Source: oshpark.com

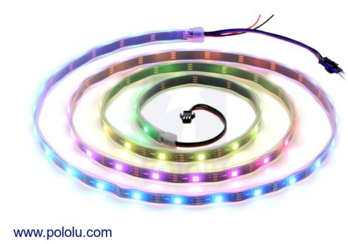

(c) Pololu LED Strip
Source: pololu.com

Figure 3.2: LED Hardware

[10] https://www.pololu.com/product/2547
[11] https://www.adafruit.com/product/2590
[12] https://oshpark.com/shared_projects/z0UBER8B
[13] http://fritzing.org/home/
[14] https://oshpark.com/

3.2.2 Software

To govern the LED strips with the micro-controller we implemented a serial communication program to flash on the metro mini. The program follows a simple design and handles four commands, a flush command which writes colors to the LED strip, a set command which sets the colors to be written to the LED strip, a clear command which clears the whole LED strip, and a command which sets the number of LEDs we have to control (Default: 60). This program is executed at the power up of the metro mini micro-controller.

To handle the communications between our BWIBot system and the micro-controller program we created an interface library, which provides six functions that can be used to control the LED strips by communicating with the micro-controller over USB serial. The six functions provided are as follows: connect, setLEDCount, clear, flush, setRGB, and setHSV. The setLEDCount and connect functions are used to initialize a serial connection with the micro-controller and to specify the number of LEDs we are controlling. The setRGB and setHSV functions both allow for the colors we desire to be set for write using RGB and HSV (Hue, Saturation, Value) formats respectively. The setHSV functions takes colors specified in HSV format and converts them to the RGB format expected by the LED strip. The flush function is used to write the set colors to the LED strip and the clear function is used to clear all colors on the LED strip.

The main control software that communicates with the rest of our software stack is written in the ROS framework.[15] The software provides some service call functions for the configuration of the LED strip parameters and testing the function

[15] All the code for the led control server is hosted in the segbot_led package under the BWI Segbot repositoryhttps://github.com/utexas-bwi/segbot/tree/master/segbot_led

of the LED strip (Figure 3.3). It also implements an action server for executing our predefined LED animations. The robot state information that we are interested in representing is related to the robot's navigational state, as such we defined seven prototype LED animations to represent important navigational states the robot can enter. These six states are as follows: left turn, right turn, blocked path, taking the elevator up, taking the elevator down, and requiring human assistance to progress in a task (Figure 3.4). In this work we used only the left and right turn animations and test the LEDs ability to communicate the robot's navigational intent.

The turn signal animations represent the state of turning using a set yellow LEDs, which cover a full beam and blink at a rate of 2 Hz (on 0.4 secs, off 0.1 secs) on the side the robot is planning to turn towards (Figure 3.4a and Figure 3.4b). The blocked animation represents the state of having the path blocked using a high intensity red pulsing animation utilizing all available LEDs simultaneously (Figure 3.4c). The elevator directional animations represent the state of wanting to take the elevator using four sets of three blue LEDs, which travel along the portions of the LED strips on each pillar of our robot simultaneously in the direction the robot is planning to go in the elevator (Figure 3.4d and Figure 3.4e). The assist animation represents the state of requiring human assistance using a low intensity blue pulsing animation similar to the blocked animation (Figure 3.4f).

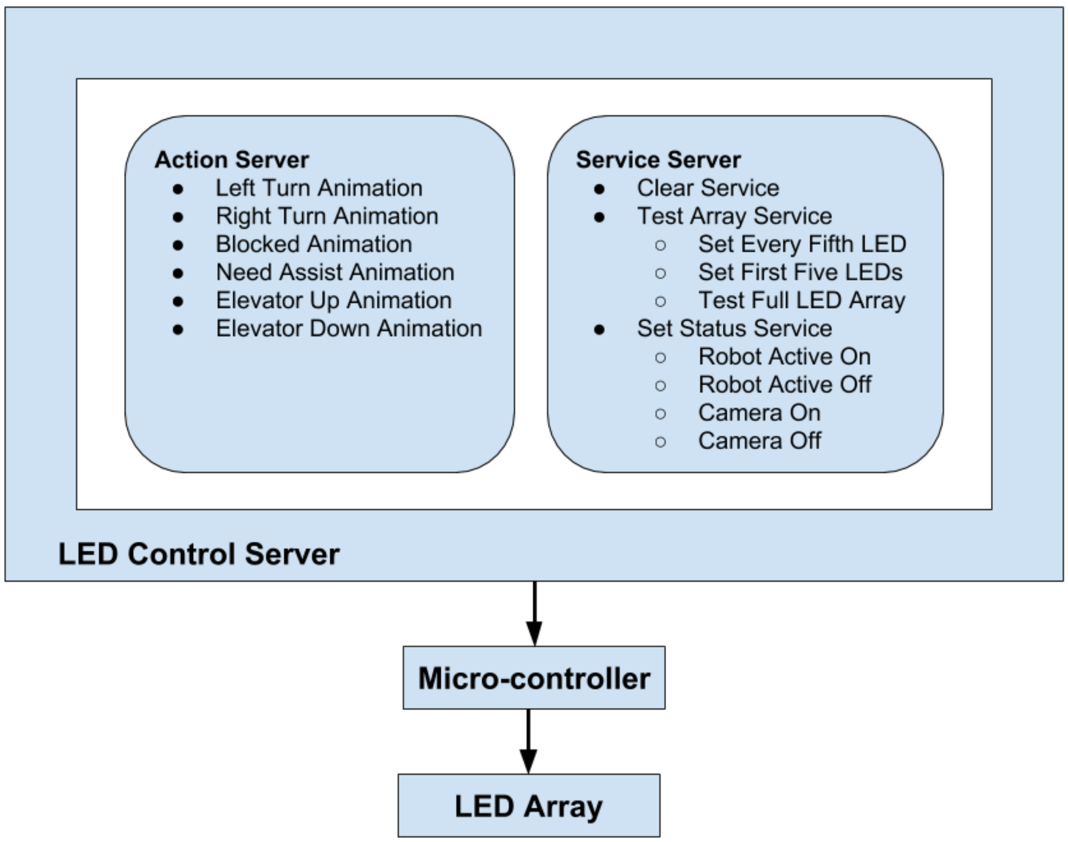

Figure 3.3: Overview of the LED Control Server structure and hierarchy of communication between the different components

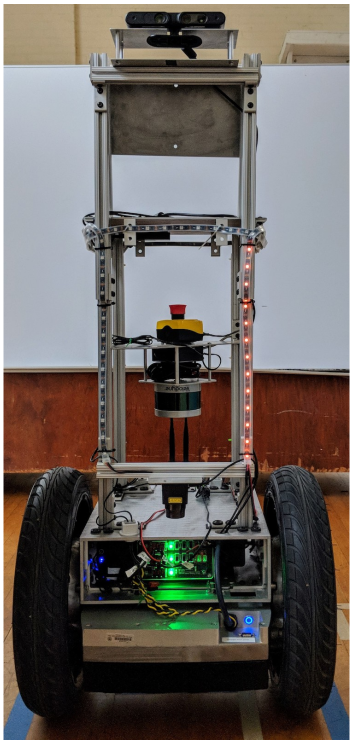
(a) Left Turn Animation


(b) Right Turn Animation


(c) Blocked Animation

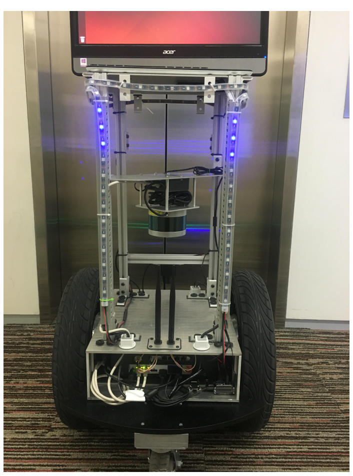
(d) Up Elevator Animation

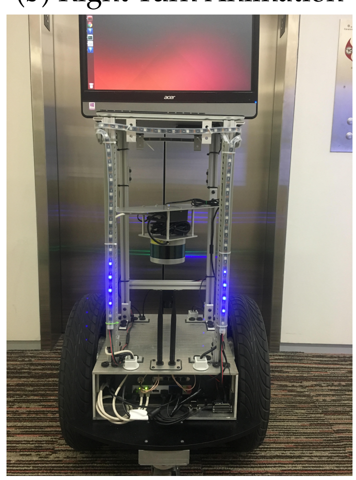
(e) Down Elevator Animation


(f) Assist Animation

Figure 3.4: LED Animations

4

Evaluation of Light-Based Signaling System

In this chapter, we focus on evaluating the viability of lights for conveying information about a mobile service robot to co-located humans. We begin by describing the setup of the experimental system and user study where we use the LED-signal to communicate intent. We then present a detailed analysis of the results of the user study. We conclude the chapter with a follow-up study on an additional condition where we use the LED-signal to communicate an instruction rather than intent.

4.1 Experimental Setup

The design of this system is based on two basic behaviors: a navigational behavior and a signaling behavior. The navigational behavior is intended to shift lanes, much as a car might in a roadway, in order to navigate around any person that it encounters who is also navigating the corridor. The signaling behavior is designed to signal the robot's intention to shift lanes to co-located humans, allowing them to adjust their own behavior in order to minimize conflict. In addition to these two behaviors, the robot may perform a brief, passive demonstration of its lane-changing behavior. This demonstration is accomplished by having the robot change lanes early in the interaction, from left to center at the opposite end of the corridor, at a distance sufficiently far away from the person that the demonstration concludes entirely before the human is close enough for the robot to begin either its signaling or lane-changing behaviors.

To perform our experiment, a corridor was constructed from cubicle furniture in a large lab space, shown in Figure 1.2. A human study participant stands at one end of the corridor, with the robot positioned at the other end. The human and robot are both instructed to traverse the hallway to the opposite end. The corridor is 17.5 meters long and 1.85 meters wide, with cameras mounted at both ends, facing down the corridor, and a third camera mounted halfway down the corridor, facing the side that the study participant begins on. The width of a hallway is based on that of the hallways in our building and its length is constrained by that of our lab space. This setup is used in a 2×2 study, where two behaviors are varied. The robot either uses its LEDs when it turns or does not and either provides a brief passive demonstration of this behavior before coming into a range in which it may come into conflict with the human or does not.

4.1.1 Navigation

For the navigational behavior, the robot splits the hallway into three lanes, as one might divide a roadway. This formulation is diagrammed in Figure 4.1. The hallway itself is 17.5 meters long and 1.85 meters wide. Each lane is 0.65 meters wide, thus there is an overlap between the lanes. If the robot detects itself to be within 1 meter of a person who is also navigating the hallway, it will stop entirely in order to allow the person to safely pass. Thus, with the width of the hallway being 1.85 meters, it is necessary for the robot and the person to be in opposite lanes, outside of the middle lane, in order to pass each other without conflict.

Three distances are defined in our model of this problem, as can be seen in Figure 4.1. Distance d_{signal}, which is at 7 meters from the person, is the distance at which the robot will signal its intention to change lanes, and is based on the

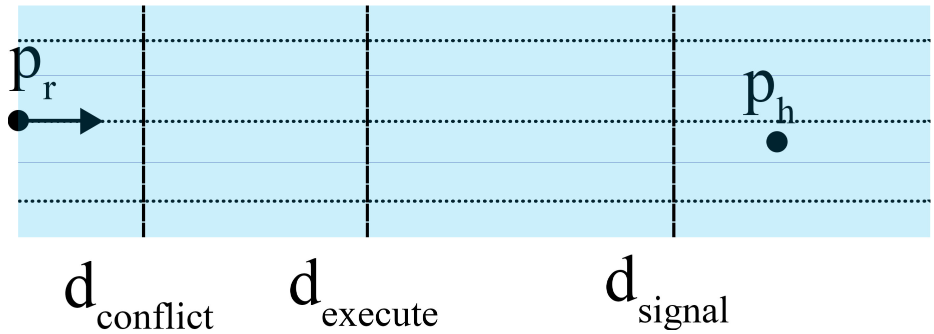

Figure 4.1: A diagram of the corridor, its lanes, and the distance thresholds at which the robot will signal its intention to change lanes, execute a lane change, and conflict with a person in its path. The distances are fixed and maintain an origin based on the position, p_r, of the robot. The respective action for each distance is executed as the position, p_h, of the person crosses the thresholds.

distance at which the robot can accurately detect a person in the hallway with its on-board sensors. The signal occurs on the side of the robot coinciding with the lane it intends to move into (Figure 3.4a and 3.4b). Distance $d_{execute}$ is the distance at which the robot will execute its turn, at 2.75 meters from the person, chosen through testing as the last possible distance to execute a turn; where choosing the same lane will ensure that a conflict will occur and choosing opposite lanes will prevent a conflict. Distance $d_{conflict}$ is the distance at which the robot determines its motion to be in conflict with that of the person and comes to a complete stop, at 1 meter from the person; chosen empirically as the minimum possible distance required to stop safely.

Because this study tests both the LEDs as a signal and the method of the passive demonstration for disambiguating this signal to the user, $d_{execute}$ is chosen to be the last possible distance that a person could choose to turn without coming into conflict with the robot. Naturally, $d_{execute}$ relies on a number of variables including the walking speed and reaction time of the person. This distance was chosen empirically - through testing the system with our lab members - but very effectively. The results demonstrate that when the robot does not signal its intention to turn, only turning when it arrives at this distance, that it comes into conflict with the person 100% of the time, as can be seen in Figure 4.2.

The navigational software is implemented as a custom ROS (30) navigation stack which attempts to minimize the distance of a point 1 meter in front of the robot to the center of the desired lane while maintaining a constant linear velocity of 0.75 m/sec. Detecting the position of the person in the corridor with respect to the robot is accomplished with a classifier that detects the person's legs in the LiDAR scan data (23). The system implements obstacle avoidance as a safeguard against possible navigational issues that could lead to a collision with a wall or the person during the course of the study, with distance ranges of $0 - 0.65$ meters for the wall and $0 - 1$ meters for the person, respectively.

4.1.2 Study Design

The study is set up as a 2×2 between-participants design with 4 conditions as in Table 1.1. Participants interacted with a robot which passively demonstrated a light change or did not, on one axis of this table; and which used an LED signal or did not, on the other[1]. The study lasted for approximately ten to fifteen minutes and consisted of four phases: (1) introduction, (2) interaction, (3) survey, (4) conclusion.

First, the experimenter explained the participant's role, obtained informed consent, and walked the participant to their starting position. Participants were instructed simply to walk to the opposite end of the corridor. No other instructions

[1] A video displaying examples of each of these conditions can be found at https://youtu.be/T4CZcP8LKRM

were given to the participants.

In phases 2 and 3, the participants traversed the corridor, and the robot did as well, creating the potential for a conflict as their paths crossed. After crossing the hallway once in this scenario, participants were administered a brief post-interaction survey comprising 35 8-point Likert and cognitive-differences scale questions describing the robot and the interaction.

In the final phase, the experimenter debriefed the participants and thanked them for their time. At this point the participants were allowed to ask any questions they might have about the study. The participants were instructed to not discuss the study with anyone as they could be potential participants for our study.

4.1.3 Participants

We recruited a total of 47 participants from The University of Texas at Austin community, 39 male, 8 female, ranging in age from 18 to 38 years. The data from 7 participants were discarded: 4 for failure to properly participate in the study (these participants stopped in front of the robot and attempted to test its capabilities or elicit responses from it), 3 due to software failures. The final pool of participants has 10 participants in each study condition.

4.2 Results

As a behavioral metric, we measure how often the robot and human study participant crossing each other's paths results in a conflict. Conflicts are defined as either the robot and the human coming to a complete stop because they come too close to one another without making a decision; the robot and the study participant entering into the same lane, forcing them to come to a stop; or scenarios in which the participant makes a rapid correction to attempt to avoid the robot, such as entering into the same lane as the robot and then changing lanes to the opposite lane, regardless of whether this causes either party to come to a stop. In the analysis of this study, conflicts were annotated based on video recorded during the interaction.

Prior to this study we conducted a pilot study of 13 participants (9 male, 4 female). It revealed a bias for participants to enter into the right lane (left lane from the perspective of the robot) in order to deconflict their path from that of the robot. Because the primary behavioral metric of this study is based on comparison of conflict scenarios, the study is designed to maximize the occurrence of these conflicts. As such, the robot always shifts into the left lane; where the participant is most likely to go by default. In addition to always going left, the robot makes its lane-change decision at the last possible moment, based on its distance from the person (2.75 meters, empirically tuned by the authors interacting with the robot). As a result, if the person has already chosen the same lane as the robot, it will almost surely result in a conflict. The purpose of this design is to maximize the impact of the intervention of introducing both the LED signal and the passive demonstration.

Results regarding the number of conflicting paths between the study participant and the robot are shown in Figure 4.2. A one-way ANOVA showed a significant main effect based on the the conditions of the study $(F(3, 36) = 9.913, p < 0.001)$. All pairwise post-hoc tests are based on Least Squares Difference (LSD), and are summarized in Table 4.1.

Results indicate almost no value to the use of the LED alone to signal the

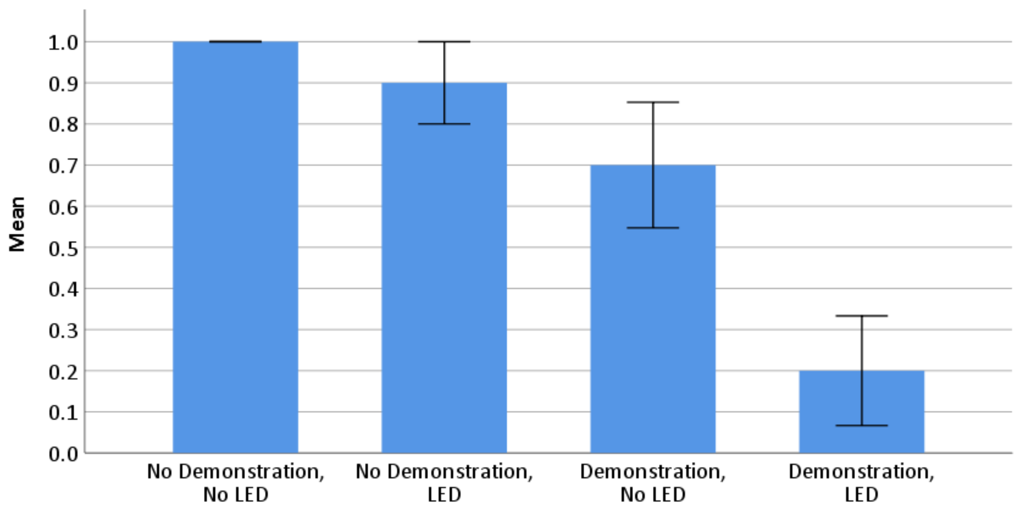

Figure 4.2: Mean Conflict Across Study Conditions.

No Demo, No LED	No Demo, LED	$MD = 0.1, p > 0.5$
No Demo, No LED	Demo, No LED	$MD = 0.3, p = 0.07$
No Demo, No LED	Demo, LED	$MD = 0.8, p < 0.01$
No Demo, LED	Demo, No LED	$MD = 0.2, p = 0.22$
No Demo, LED	Demo, LED	$MD = 0.7, p < 0.01$
Demo, No LED	Demo, LED	$MD = 0.5, p < 0.01$

Table 4.1: Pairwise Comparison with LSD Post-Hoc Test.

robot's turning behavior. Post-hoc tests showed no significant difference between showing the LED with no passive demonstration (the No Demonstration, LED condition) and simply not using the LED (No Demonstration, No LED), with only a modest mean difference ($MD = 0.1, p > 0.5$). However, the story changes once the passive demonstration technique is introduced. Passively demonstrating the robot's lane-shifting behavior - having the robot perform a lane-shift in front of the participant further down the hallway, prior to their immediately proximal interaction with potential for conflict - is sufficient to reduce conflicts to a nearly-significant level (No Demonstration, No LED vs Demonstration, No LED:

$MD = 0.3, p = 0.07$). Moreover, when combined with the LED and despite the ineffectiveness of the LED signal on its own, this effect is compounded to a large margin at a significant level (No Demonstration, No LED vs Demonstration, LED: $MD = 0.8, p < 0.01$). It can be seen that the compound effect of demonstrating the signal, rather than simply demonstrating the lane-shift, is what makes this effect so powerful (Demonstration, No LED vs Demonstration, LED: $MD = 0.5, p < 0.01$).

The strength of the passive demonstration with the LED signal is further reflected in the results of the post-interaction survey, shown in Figure 4.3. None of the questions asked had a significant main effect, though here we present the responses for a few interesting questions. Participants appeared to interpret the robot's LED signal as an attempt to communicate (Figure 4.3a) and indicated that communication was most clear when the robot performed a passive demonstration alongside the LED signal (4.3b). The survey responses support the idea that the LED signal is not very useful on its own, while further reinforcing the value of the demonstration. Furthermore, a passive demonstration in the absence of a signal appears to have harmed performance on this metric, which may be reflective of participants noting a lack of communicative signaling during lane-shifting after witnessing the behavior twice with no communicative signal; despite the reduction in conflict in both passive demonstration conditions.

It is unsurprising that the interaction was found to be most comfortable in the presence of the passive training and the LED, Figure 4.3c, though the difference between "No Demonstration, No LED" and "Demonstration, LED" is only quite small, with the worst performing conditions being "No Demonstration, LED" and "Demonstration, No LED." This result could either indicate that a larger sample size is necessary to clarify this response, or potentially that the LED or demonstra-

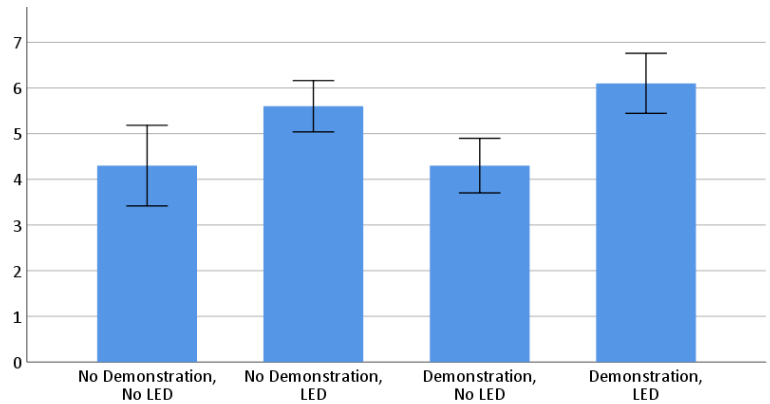

(a) “The robot was trying to communicate.”

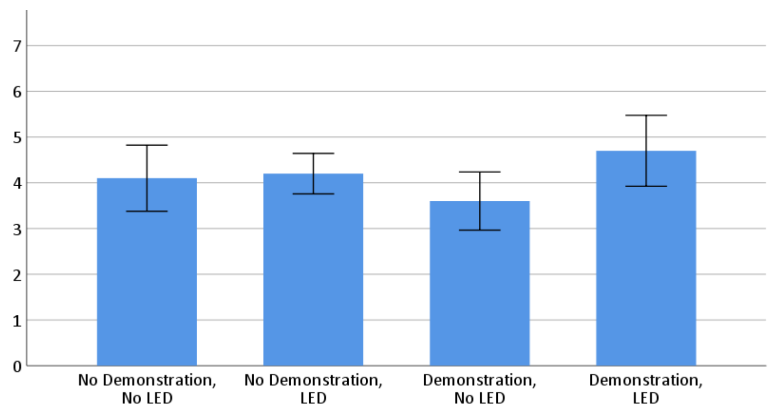

(b) “The robot communicated its intentions clearly.”

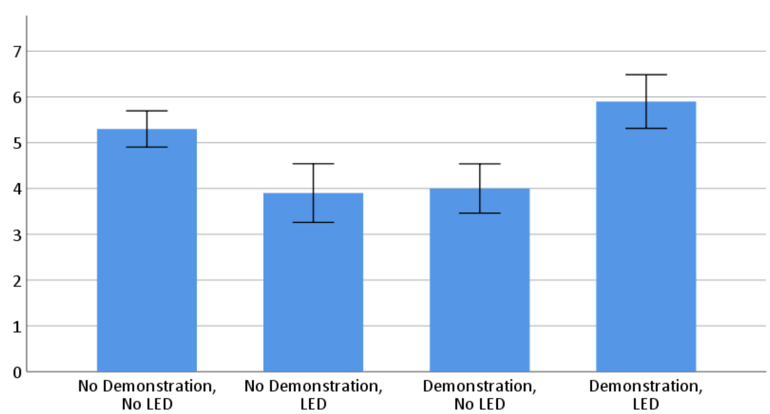

(c) “The interaction was comfortable.”

Figure 4.3: Responses to Likert-scale questions from the post-interaction survey.

tion on their own are simply confusing to participants. This also points toward a potentially-confounding variable and limitation in the design of this study, which is that the overall path of the robot in the conditions in which it provides a passive demonstration is curvier, as it must make two lane-changes instead of one in order to provide the demonstration. Eliminating this issue will require observing this behavior in additional scenarios so as to determine the magnitude of the overall effect of the curviness of the path. This may, indeed, be reflected in our survey results for the questions, "The robot communicated its intentions clearly" (Figure 4.3b) and "The interaction was comfortable" (Figure 4.3c).

4.3 Follow-up Study

In our pilot study, participants indicated roughly half of the time that they thought that the LED signal was an instruction regarding which lane they should go into, rather than an indication of the robot's intended heading. To address these responses an additional LED signal condition was evaluated, where the signal communicated an instruction to the participant instead of the robot's intent without the passive demonstration. The results of this condition demonstrated a similar modest improvement in reducing the number of navigational conflicts as in the case of using the LED signal for communicating intent without the passive demonstration.

4.3.1 Participants

We recruited an additional of 11 participants from The University of Texas at Austin community, 8 male, 3 female, ranging in age from 19 to 31 years. The data of 1 participant was discarded due to software failures. The final pool of participants has a total 10 participants in for the additional condition.

4.3.2 Measures and Analysis

For these additional results we use the same measures as those introduced in Section 4.2.

The results of using the signal for communicating and instruction without the passive demonstration demonstrated a similar modest improvement in reducing the number of navigational conflicts as in the case of using the LED signal for communicating intent without the passive demonstration. Using the signal for instruction resulted in conflict 8 of 10 times, where the using the signal for communicating intent resulted in conflict 9 out of 10 times (Figure 4.4). A one-way ANOVA showed a significant main effect based on all the conditions of the study ($F(4, 45) = 7.040, p < 0.001$). All pairwise post-hoc tests are based on Least Squares Difference (LSD), and are summarized in Table 4.2.

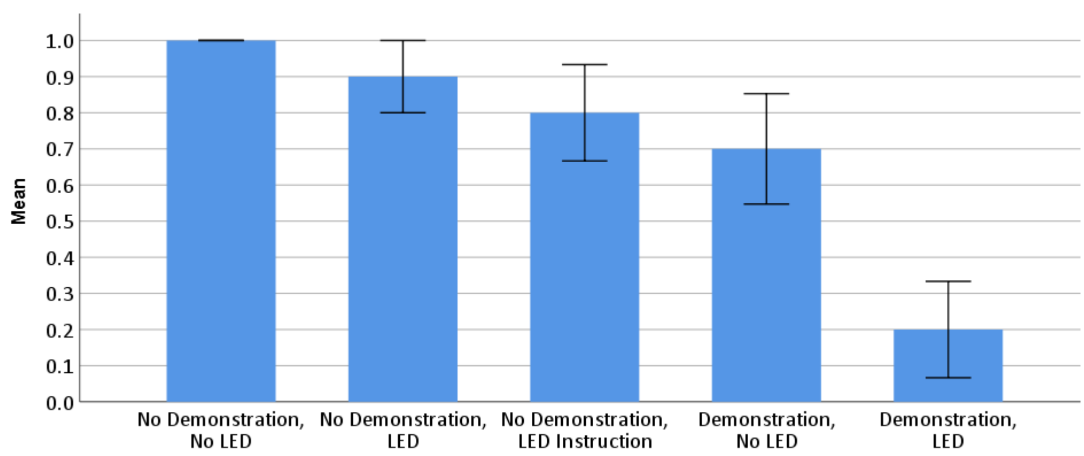

Figure 4.4: Mean Conflict Across Study Conditions with Additional LED Instruction Condition

No Demo, LED Instruction	No Demo, No LED	$MD = 0.2, p = 0.235$
No Demo, LED Instruction	No Demo, LED	$MD = 0.1, p > 0.5$
No Demo, LED Instruction	Demo, No LED	$MD = 0.1, p > 0.5$
No Demo, LED Instruction	Demo, LED	$MD = 0.6, p < 0.01$
No Demo, No LED	No Demo, LED	$MD = 0.1, p > 0.5$
No Demo, No LED	Demo, No LED	$MD = 0.3, p = 0.07$
No Demo, No LED	Demo, LED	$MD = 0.8, p < 0.01$
No Demo, LED	Demo, No LED	$MD = 0.2, p = 0.22$
No Demo, LED	Demo, LED	$MD = 0.7, p < 0.01$
Demo, No LED	Demo, LED	$MD = 0.5, p < 0.01$

Table 4.2: Pairwise Comparison with LSD Post-Hoc Test for LED Instruction Condition.

Results still indicate almost no value to the use of the LED alone even when they are communicating an instruction instead of intent. Post-hoc tests showed no significant difference between communicating intent using the LED with no passive demonstration (the No Demonstration, LED condition) and communicating an instruction using the LED (the No Demonstration, LED Instruction condition), with only a modest mean difference ($MD = 0.1, p > 0.5$). This is further reflected in the results of the post-interaction survey, shown in Figure 4.5. None of the questions asked had any significant difference between these two conditions, though here we present the responses for the same few interesting questions as those in Section 4.2. Participants appeared to interpret the robot's LED Signal Instruction condition as an attempt to communicate ($MD = 0.5, p > 0.5$) (Figure 4.5a), though indicated that communication was less clear when the robot was using this condition ($MD = 0.7, p = .472$) (4.5b). The survey responses still support the idea that the LED signal is still not very useful on its own even with this additional condition.

It was expected that the responses for the additional condition "No Demonstration, LED Instruction" and the previous condition "No Demonstration, LED" would be quite similar as not much changed besides the type of information being communicated. Between this additional condition and simply not using the LED (the No Demonstration, No LED condition) there was a slightly lower p-value ($MD = 0.2, p = 0.235$), where previously we had ($MD = 0.1, p > 0.5$) when communicating intent. Though the p-value is lower it still does not provide any significant effect overall. This result could either indicate that a larger sample size is necessary to clarify this response, or potentially that the LEDs on their own are simply confusing to participants regardless of whether an instruction or intent is being communicated.

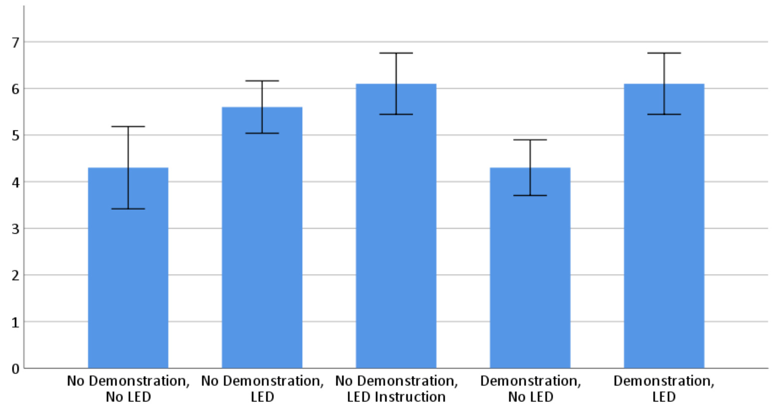

(a) "The robot was trying to communicate."

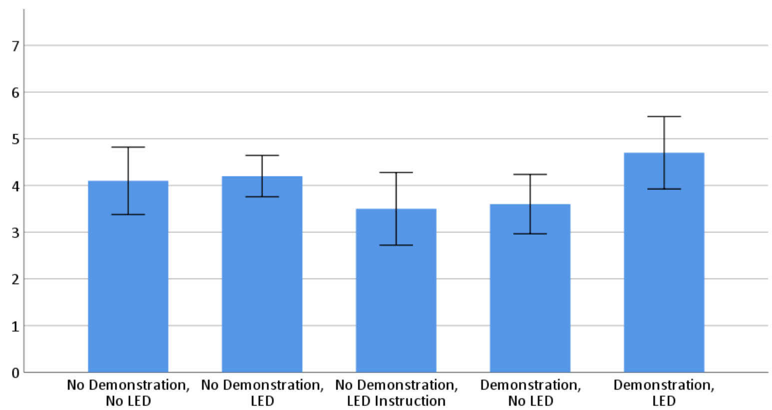

(b) "The robot communicated its intentions clearly."

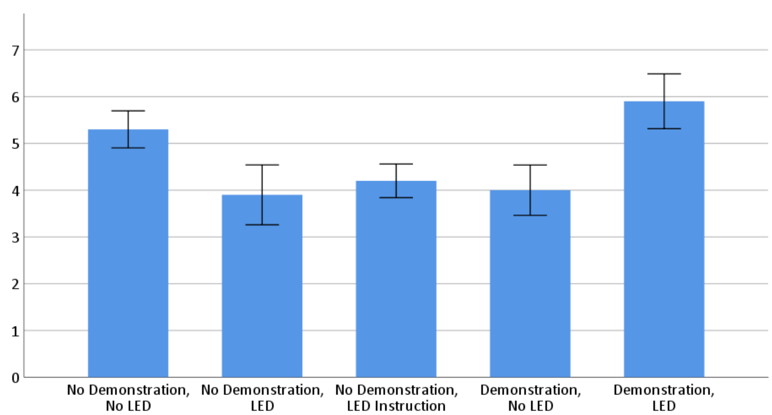

(c) "The interaction was comfortable."

Figure 4.5: Responses to Likert-scale questions from the post-interaction survey with the LED Instruction Condition.

5

Conclusion

In this thesis, we present an LED-based signal control system implemented in ROS that allows for various animations to be executed on an array of LEDs. This system is implemented on our BWIBots allowing them to communicate information to co-located humans using LED signals. The LED signals allow the robots to continually communicate information in a modality that is less intrusive then repetitive speech communication of the same message and more engaging than text messages on an on-screen display.

This thesis demonstrated that a familiar signal, the turn signal, was difficult to understand when divorced from the context of driving. However, when passive demonstrations were introduced, the meaning became clear. Our initial assumption on crafting the pilot study was that users would find turn signals on a robot to be entirely intuitive, and that the experience that they had gained from driving - which involves navigating a space shared with other cars - would directly transfer to the task of navigating a space shared with a robot. This was not the case. Instead, our study revealed that simple passive demonstrations of the signal to study participants were sufficient to disambiguate the intention of this signal to them.

These findings leave open, the question as to whether this method would easily extend to other novel signals. Is it that the robot's demonstration, combined with the familiarity of turn signals, gives rise to our result; or that this technique would extend to entirely novel signals that are not encountered in day-to-day life? Probing this question will require the design of studies which use non-anthropomorphic signals in scenarios that are not commonly encountered outside of interactions with robots, such as collaborative manipulation tasks. We hope to further explore this question in experiments with newer versions of the BWIBot which incorporate robotic arms.

This result will inform the design of future versions of our BWIBot's navigation stack, whereby we intend to add turn signals to the robot, but also to program the robot to perform a few quick lane-shift and turning maneuvers in front of unfamiliar faces; thus enabling new users to acclimate to the behavior of the system. Once complete, we will perform studies of the system involving larger crowds and longitudinal studies of the deployed system on the BWIBot platform. These studies are currently in the planning phases. Our result demonstrates a scenario in which passive demonstrations are able to provide the context required to interpret the robot's LED signal to users, and opens the study of such passive demonstrations an interesting new direction for further research.

Bibliography

[1] Kim Baraka. Effective Non-Verbal Communication for Mobile Robots using Expressive Lights. Master's thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, May 2016.

[2] Kim Baraka and Manuela M. Veloso. Mobile Service Robot State Revealing Through Expressive Lights: Formalism, Design, and Evaluation. *International Journal of Social Robotics*, 10(1):65–92, Jan 2018. ISSN 1875-4805. doi: 10.1007/s12369-017-0431-x. URL `https://doi.org/10.1007/s12369-017-0431-x`.

[3] Cindy L Bethel. *Robots Without Faces: Non-Verbal Social Human-Robot Interaction*. PhD thesis, Computer Science and Engineering, University of South Florida, 2009.

[4] Cindy L Bethel and Robin R Murphy. Survey of non-facial/non-verbal affective expressions for appearance-constrained robots. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(1):83–92, 2008.

[5] Elizabeth Cha, Yunkyung Kim, Terrence Fong, Maja J Mataric, et al. A Survey of Nonverbal Signaling Methods for Non-Humanoid Robots. *Foundations and Trends® in Robotics*, 6(4):211–323, 2018.

[6] Steve Cousins. Ros on the pr2 [ros topics]. *IEEE Robotics & Automation Magazine*, 17(3):23–25, 2010.

[7] Winand H Dittrich and Stephen EG Lea. Visual perception of intentional motion. *Perception*, 23(3):253–268, 1994.

[8] Winand H Dittrich, Tom Troscianko, Stephen EG Lea, and Dawn Morgan. Perception of emotion from dynamic point-light displays represented in dance. *Perception*, 25(6):727–738, 1996.

[9] Kotaro Funakoshi, Kazuki Kobayashi, Mikio Nakano, Seiji Yamada, Yasuhiko Kitamura, and Hiroshi Tsujino. Smoothing human-robot speech interactions by using a blinking-light as subtle expression. In *Proceedings of the 10th international conference on Multimodal interfaces*, pages 293–296. ACM, 2008.

[10] Martin Gebser, Roland Kaminski, Benjamin Kaufmann, and Torsten Schaub. Clingo = ASP+ control: Preliminary report. *arXiv preprint arXiv:1406.3694*, 2014.

[11] Rebecca Hansson and Peter Ljungstrand. The reminder bracelet: subtle notification cues for mobile devices. In *CHI'00 Extended Abstracts on Human Factors in Computing Systems*, pages 323–324. ACM, 2000.

[12] Chris Harrison, John Horstman, Gary Hsieh, and Scott Hudson. Unlocking the expressivity of point lights. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1683–1692. ACM, 2012.

[13] Kunimatsu Hashimoto, Fuminori Saito, Takashi Yamamoto, and Koichi Ikeda. A field study of the human support robot in the home environment. In *Advanced Robotics and its Social Impacts (ARSO), 2013 IEEE Workshop on*, pages 143–150. IEEE, 2013.

[14] Jesse Thomason, Aishwarya Padmakumar, Jivko Sinapov, Justin Hart, Peter Stone, and Raymond J. Mooney. Opportunistic Active Learning for Grounding Natural Language Descriptions. In SSergey Levine, Vincent Vanhoucke, and Ken Goldberg, editors, *Proceedings of the 1st Annual Conference on Robot earning (CoRL-17)*, pages 67–76, Mountain View, California, November 2017. PMLR. URL http://www.cs.utexas.edu/users/ai-lab/pub-view.php?PubID=127657.

[15] Piyush Khandelwal and Peter Stone. Multi-robot human guidance using topological graphs. In *AAAI Spring 2014 Symposium on Qualitative Representations for Robots*, 2014.

[16] Piyush Khandelwal, Fangkai Yang, Matteo Leonetti, Vladimir Lifschitz, and Peter Stone. Planning in action language BC while learning action costs for mobile robots. In *Proceedings of the 24th International Conference on Automated Planning and Scheduling*, pages 472–480. Association for the Advancement of Artificial Intelligence (AAAI), 2014.

[17] Piyush Khandelwal, Samuel Barrett, and Peter Stone. Leading the Way: An Efficient Multi-robot Guidance System. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1625–1633. AAMAS, 2015.

[18] Min-gyu Kim, Hui Sung Lee, Jeong Woo Park, Su Hun Jo, and Myung Jin Chung. Determining color and blinking to support facial expression of a robot for conveying emotional intensity. In *RO-MAN 2008-The 17th IEEE International Symposium on Robot and Human Interactive Communication*, pages 219–224. IEEE, 2008.

[19] Kazuki Kobayashi, Kotaro Funakoshi, Seiji Yamada, Mikio Nakano, Takanori Komatsu, and Yasunori Saito. Blinking light patterns as artificial subtle expressions in human-robot speech interaction. In *RO-MAN 2011-The 20th IEEE International Symposium on Robot and Human Interactive Communication*, pages 181–186. IEEE, 2011.

[20] Scott R Kuindersma, Edward Hannigan, Dirk Ruiken, and Roderic A Grupen. Dexterous mobility with the uBot-5 mobile manipulator. In *International Conference on Advanced Robotics*, pages 1–7. IEEE, 2009.

[21] Irving Langmuir and WF Westendorp. A study of light signals in aviation and navigation. *Journal of Applied Physics*, 1(5):273–317, 1931.

[22] Patricio Lankenau. Virtour: Telepresence System for Remotely Operated Building Tours. Turing thesis, Department of Computer Science, The University of Texas at Austin, Austin, TX, August 2016.

[23] Angus Leigh, Joelle Pineau, Nicolas Olmedo, and Hong Zhang. Person Tracking and Following with 2D Laser Scanners. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 726–733. IEEE, 2015.

[24] Vladimir Lifschitz. What Is Answer Set Programming?. In *AAAI*, volume 8, pages 1594–1597, 2008.

[25] James E Lloyd. Bioluminescence and communication in insects. *Annual review of entomology*, 28(1):131–160, 1983.

[26] Andrii Matviienko, Andreas Löcken, Abdallah El Ali, Wilko Heuten, and Susanne Boll. Navilight: investigating ambient light displays for turn-by-turn navigation in cars. In *Proceedings of the 18th International Conference on Human-Computer Interaction with Mobile Devices and Services*, pages 283–294. ACM, 2016.

[27] Samuel Finley Breese Morse. *Electro-magnetic Telegraph*. US Government Printing Office, 1844.

[28] Piyush Khandelwal and Peter Stone. Multi-Robot Human Guidance: Human Experiments and Multiple Concurrent Requests. In *International Conference on Autonomous Agents and Multiagent Systems*. AAMAS, May 2017.

[29] Piyush Khandelwal, Shiqi Zhang, Jivko Sinapov, Matteo Leonetti, Jesse Thomason, Fangkai Yang, Ilaria Gori, Maxwell Svetlik, Priyanka Khante, Vladimir Lifschitz, J. K. Aggarwal, Raymond Mooney, and Peter Stone. BWIBots: A platform for bridging the gap between AI and human–robot interaction research. *The International Journal of Robotics Research*, 2017. doi: 10.1177/0278364916688949.

[30] 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 *ICRA workshop on open source software*, volume 3, page 5. Kobe, Japan, 2009.

[31] Nicolaus A Radford, Philip Strawser, Kimberly Hambuchen, Joshua S Mehling, William K Verdeyen, A Stuart Donnan, James Holley, Jairo Sanchez, Vienny Nguyen, Lyndon Bridgwater, et al. Valkyrie: NASA's first bipedal humanoid robot. *Journal of Field Robotics*, 32(3):397–419, 2015.

[32] Ulrich Reiser, Christian Connette, Jan Fischer, Jens Kubacki, Alexander Bubeck, Florian Weisshardt, Theo Jacobs, Christopher Parlitz, Martin Hägele, and Alexander Verl. Care-O-bot® 3-creating a product vision for service robot applications by integrating design and technology. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 1992–1998. IEEE, 2009.

[33] Stephanie Rosenthal, Joydeep Biswas, and Manuela Veloso. An effective personal mobile robot agent through symbiotic human-robot interaction. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pages 915–922. AAMAS, 2010.

[34] Shiqi Zhang, Yuqian Jiang, Guni Sharon, and Peter Stone. Multirobot Symbolic Planning under Temporal Uncertainty. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Sytems*. AAMAS, May 2017.

[35] Jivko Sinapov, Priyanka Khante, Maxwell Svetlik, and Peter Stone. Learning to Order Objects Using Haptic and Proprioceptive Exploratory Behaviors. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. IJCAI, 2016.

[36] Siddhartha S Srinivasa, Dmitry Berenson, Maya Cakmak, Alvaro Collet, Mehmet R Dogar, Anca D Dragan, Ross A Knepper, Tim Niemueller, Kyle Strabala, Mike Vande Weghe, et al. Herb 2.0: Lessons learned from developing a mobile manipulator for the home. *Proceedings of the IEEE*, 100(8): 2410–2428, 2012.

[37] Matthew Sweeney, Thomas Pilarski, William Payne Ross, and Chenggang Liu. Light Output System for a Self-Driving Vehicle, Mar 15 2018. US Patent App. 15/803,184.

[38] Daniel Szafir, Bilge Mutlu, and Terry Fong. Communicating directionality in flying robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*, pages 19–26. ACM, 2015.

[39] Jesse Thomason, Shiqi Zhang, Raymond Mooney, and Peter Stone. Learning to interpret natural language commands through human-robot dialog. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*. IJCAI, 2015.

[40] Jesse Thomason, Jivko Sinapov, Maxwell Svetlik, Peter Stone, and Raymond J Mooney. Learning Multi-Modal Grounded Linguistic Semantics by Playing" I Spy. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. IJCAI, 2016.

[41] Manuela M Veloso, Joydeep Biswas, Brian Coltin, and Stephanie Rosenthal. CoBots: Robust Symbiotic Autonomous Mobile Service Robots. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, pages 4423–4429. IJCAI, 2015.

[42] Fangkai Yang, Piyush Khandelwal, Matteo Leonetti, and Peter Stone. Planning in answer set programming while learning action costs for mobile robots. In *AAAI Spring 2014 Symposium on Knowledge Representation and Reasoning in Robotics*. AAAI-SSS, 2014.

[43] Shiqi Zhang, Mohan Sridharan, Michael Gelfond, and Jeremy Wyatt. Towards an architecture for knowledge representation and reasoning in robotics. In *International Conference on Social Robotics*, pages 400–410. Springer, 2014.

[44] Shiqi Zhang, Dongcai Lu, Xiaoping Chen, and Peter Stone. Robot Scavenger Hunt: A Standardized Framework for Evaluating Intelligent Mobile Robots. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. IJCAI, 2016.