Leading a Visitor Intelligently

Group Members: Luke Wright, Stone Tejeda, Katherine Hackworth

Application Task/Problem

Currently the robots can successfully navigate to a given goal, and recognize people with some level of reliability. However, functionality to lead a person to a given objective is limited. The goal we would like to accomplish is to implement functionality to lead an individual to that given objective, with a focus on Human-Robot Interaction (HRI) to lead the person intelligently and reliably. The robot will be given a goal by the person either through a command, or perhaps in future iterations a voice command, and lead that person to their goal, ideally communicating with them and periodically checking in to make sure it has not lost the person.

Previous Experience

As a group, we do not have much experience in this problem beyond what we have done within this class; however, prior homeworks do have some relevance to the task we will be working on. Our work with basic planning will serve as a starting point for how we will implement the robot's declaration of its movements for persons to be aware of where the robot is going. Our work with computer vision will also serve as a starting point with determining whether or not a person is still following the robot. While our initial goal will be to implement a simple text interface to determine if someone is still following the robot, we intend to move to using the rear-facing camera on the v3 BWIBot to check for a follower if we have accomplished our initial goals. Due to there only being a rear-facing camera on the v3 BWIBot, that portion of the project would have to be implemented only on that platform.

Expected Achievements

We have categorized our goals into 5 subgroups of achievability. The first group of goals are those which we absolutely intend to have done. These absolute goals will be having a command interface for selecting navigation goals which does not interfere with the existing goal-sending interface, and creating a text confirmation system which asks the user whether he or she is still following the robot.

Once those goals are met, we intend to move on to the second subgroup of goals, which we believe will likely be accomplished before the semester ends, barring some unforeseen issues within the first subgroup of goals. Within this subgroup are two goals: using the existing GUI and LEDs to announce turns, and having the robot choose

to return to the lab if it has determined that there is no longer a person following it. Neither of these goals seem to be particularly challenging, so we expect to have these done.

Within the third subgroup is one goal, which we do not fully expect to have completed by the end of the semester. This is to have the robot announce turns and room entries via text-to-speech. Due to this being a similar goal to announcing turns via a GUI and the LEDs, we believe that accomplishing this goal is feasible, but will be satisfied with the level of accomplishment we have achieved if we do not make it to this subgroup of goals.

We do not expect to accomplish the fourth subgroup, but believe that these goals could potentially be done in the future, or perhaps within the deadline if we end up ahead of schedule. These goals are using the rear-facing camera to determine if the user is still following the robot, and using speech-to-text to create the navigation goal. We believe that both of these goals present far more challenging issues due to the usage of computer vision and natural language processing.

The final subgroup is a goal which we do not expect to have accomplished by the end of the semester, but believe is worth mentioning regardless, as it would present interesting information from an HRI perspective and is relevant to the project. This is to have the robot maintain a specified distance from the user using the rear-facing camera. This would require real-time monitoring of the user, which would become very complicated in busy areas and could prove to take up too much processing time to even be feasible. Nevertheless, data gathered from this goal would be very interesting as it would show how people measure their personal space when interacting with a robot.

Evaluating Success

Determining success with our project can be broken down into a number of subgroups. Firstly, the most simple way to evaluate success would simply be to ensure that the robot is still successfully reaching the requested destination. While this goal seems easily achievable, it is worth mentioning simply because there may be a chance that our own code interferes with the planning and goal-reaching algorithms already in place, and we want our project to coexist with that existing functionality, rather than interfere with it.

Our second point of evaluation will be centered around how frequently the robot manages to successfully guide a person from their current location to the goal. Due to there being a possibility of persons in the real world simply abandoning the robot, we will measure success not on whether the robot makes it to the goal and still has a person with them, but whether the robot is aware of there being a person with it once it has reached the goal. Using this tool, we hope to have the robot accurately know whether it is still being followed 80% of the time. At first, this will be tested by having a

single person follow the robot around in an area with low traffic, and the follower will consistently respond to the robot's requests to confirm that it is still being followed. Beyond that, we hope to have the user occasionally fail to respond to the robot, and also place the robot in more crowded areas. This second evaluation would not necessarily be more difficult when we are using the text confirmation to determine if the robot is still being followed, but will likely prove to be more difficult should we move to having the robot confirm that it is being followed through the rear-facing camera.

Our last form of evaluation is not actually a measure of success, but an evaluation of information from an HRI perspective. We intend to have the robot survey participants about their opinions on how it interacted with them. The questions we intend to ask are as follows:

- Whether the user felt as though the robot asked for following confirmation too frequently or infrequently
- Whether the robot was too close or too far away from the user (this question will only be asked if we are using the rear-facing camera to ensure that the robot is a certain distance away from the user)

We intend to vary the confirmation frequency and distance from user throughout the experiment, and examine what distance and frequency seem to garner the best interaction between robot and user.

Related Work

There have been several past projects done on the BWI bots which relate and should provide working bases for our proposed project. First, the bots can currently plan routes to a given objective which would likely be outside the scope of our project [1]. Additionally, there has been work done in prior semesters with the LED lights being used as turn signals, so there should presumably be some existing methods of determining what constitutes a "turn" for the robot, we hope to extend this functionality to inform the visitor of upcoming actions so they can follow more easily [2]. The bots also have some level of text-to-speech functionality currently, which will allow us to address the visitor not only through a graphical user interface (GUI) but also through voice, which will be useful as the individual could likely become distracted from the screen in a busy building environment. Finally, the robots currently have basic functionality to recognize a person visually [3] and the V3 has a rear-facing camera, this functionality will be useful in determining whether the individual has been lost or not. We hope to either be able to interface with those projects through ROS's messaging functionality so that we do not have to reinvent the functionality they provide.

Projected timeline of tasks/milestones

April 3rd	Turn in proposal, freeze goals for project, make github for project
April 9th	Complete surveying existing functionality and how to interface with it via ROS, have rough outline of implementation approach
April 23rd	Command to lead to location works (robot goes to location), text confirmation
April 30th	GUI + lights for leading functionality, graceful failure
May 6th	Data collection, tweaks
May 11th	Turn in final report

References:

1: https://github.com/utexas-bwi/segbot_apps
2: http://www.cs.utexas.edu/~jsinapov/teaching/cs378/reports/wrench_phan.pdf
3: http://www.cs.utexas.edu/~jsinapov/teaching/cs378/reports/brenan_sadami_zhou.pdf