

The Bryn Mawr Tour Guide Robot

Catherine Chiu

In partial fulfillment
for the Senior Thesis requirement
toward an independent major in Computer Science.

April 26, 2004

Abstract

In this thesis I will discuss how a robotic tour guide that travels and displays two wings of Bryn Mawr College's Park Science Building is created. The robot gives a pre-scripted, pre-planned tour using dynamic navigation and localization. Localization is carried out implementing a combination of dead reckoning and landmark detection which is based on vision. The robot interacts with visitors using speech recognition which has a limited vocabulary for recognition. The robot also uses speech to narrate the pre-scripted tour and give directions to visitors. I will also provide an overview of previous tour-giving robots as well as examine control strategies and details of our implementation. While most similar projects are of the Ph.D. level, our focus will be to integrate several components to create the tour guide robot.

Contents

1	Introduction	3
2	Related Work	4
2.1	Rhino	4
2.2	Minerva	6
2.3	Alfred	8
2.4	Grace	10
2.5	Virgil	12
3	Architecture of The Bryn Mawr Tour Guide Robot	12
3.1	Navigation	14
3.1.1	Obstacle Avoidance	14
3.2	Localization	15
3.3	Human Interaction	16
4	Current Status	16
4.1	Navigation	17
4.2	Localization	17
4.3	Speech (Human Interaction)	18
4.3.1	Direction Giving	18
4.3.2	Tour Giving	20
5	Conclusion	20

References

Appendix A

Appendix B

List of Figures

1	Representation of the smaller environment of Park.	3
2	Picture of Rhino.	5
3	Minerva	6
4	This is a diagram of Minerva's software architecture.	7
5	Swarthmore College's Alfred.	8
6	Description of Alfred's finite state machine.	9
7	Grace.	10
8	This is a picture of the PeopleBot robot with its indicated features.	13
9	Representation of our robot's architecture.	13
10	The robot in the hallway. Notice the landmarks on the floor.	15
11	Diagram of our finite state machine.	17
12	A depiction of the points that the directs mention.	19
13	Picture of the planned tour with the red and blue land markers in place.	21




Figure 1: Representation of the smaller environment of Park.

1 Introduction

The goal of this thesis is to explore creating a robot that gives tours of Bryn Mawr College's science building. The motivation behind this project was to create more awareness of the science departments for prospective students and their families as well as act as an aid to visitors, current students, and faculty in providing navigation. The robot will use human interaction to provide directions around the building. This project was sponsored by a research grant from The Committee on the Status of Women in Computing Research (CRA-W) for the academic year of 2003-2004. The grant was awarded to Ioana Butoi, Darby Thompson, and myself under the advisement of Professor Douglas Blank.

The Park Science Building is a large labyrinth-like building which houses several departments and their laboratories. This building is confusing even to its regular occupants and is often heavily avoided by students in subjects outside of math and the sciences. Due to the large size of the building, the environment was broken down to a manageable size consisting of the building's main entrance and the hallway that goes by the Computer Science and Physics classrooms and laboratories (see Figure 1).

Major components of this project included using Pyro to integrate navigation, localization, obstacle avoidance, and human interaction. Pyro is an application being developed by Prof. Douglas Blank, Prof. Deepak Kumar, Prof. Lisa Meeden, and others such that multiple robots have a common platform making it possible for user's to not be concerned with most low level details [2, 3].

2 Related Work

There have been other related projects using algorithms and techniques that we hope to implement and integrate. In what follows, I will present several successful robotics projects and the algorithms. I will also discuss ideas and algorithms that we considered or adapting for our tour guide. Although most of these projects were Ph.D. level, our focus was to integrate and adapt ideas from previous projects such as these described below.

2.1 Rhino

Rhino was an indoor navigation robot designed at the University of Bonn in Germany in the mid 1990s [4]. Rhino's main goals were to be able to navigate within an office space, move from point A to point B, and recognize a number of objects on the floor and dispose of them in a garbage can. Rhino (see Figure 2) used occupancy grids built from the data received from its 24 sonar sensors and two color cameras. Neural networks were used to interpret the sonar data to estimate the likelihood of occupancy by a wall or doorway within a three meter radius of the robot. Rhino also used the sonar data to locate immediate obstacles and create a trajectory to navigate around the obstructions. Rhino was considered successful in achieving these goals, but lacked the consistency that autonomous agents need.



Figure 2: Picture of Rhino.

Rhino was later adapted in 1997 with improvements and new features, such as safe navigation and an appealing user interface [5]. In collaboration with Carnegie Mellon University, Rhino was adapted to give a tour of the Deutsches Museum Bonn in Bonn, Germany for six days. This newer version of Rhino had the addition of laser sensors which provided more accurate and reliable data about its surroundings.

Since the environment was heavily populated at times and often changing, safe and accurate navigation was a challenging task. This problem was solved by using Markov localization. Markov localization is an algorithm that uses sensor data to estimate the location of the robot by assigning a probability distribution of each option, therefore eliminating options that are not likely as there is a possibility that the path is obstructed.

Rhino also included a user friendly interface which communicated with users through its small digital screen and received input from four buttons. Users could also communicate with Rhino via the Internet and receive an interactive tour with the use of Java applets.

Rhino was very successful in giving safe and accurate tours in the museum. Rhino is an example of how occupancy grids ¹ are an effective algorithm for navigation and how

¹An occupancy grid is a representation of the current environment reflecting the possibility of an obstacle. The more intense the color, the higher the probability of occupancy.



Figure 3: Minerva

important it is for an autonomous agent have a user friendly interface.

2.2 Minerva

Minerva (see Figure 3) is a second generation tour guide robot created by a team made up of members from Carnegie Mellon University and University of Bonn, several of which had worked on Rhino [11]. Minerva was featured at the Smithsonian's National Museum of American History in Washington, D.C. for two weeks in the Summer of 1998. Minerva's goal was to entertain and educate museum visitors about the exhibits by guiding them through the museum and providing brief explanations of some exhibits. Minerva used some of the same technology as Rhino, such as navigation and path planning modules, however with some improvements. Minerva used occupancy grids, like Rhino, but used a textured map² as well to ensure safety and accuracy in navigation. Both its occupancy grids are based on sensor readings from Minerva's environment. Sensors included laser scans, odometry readings³, and camera images (see Figure 3). A human controlled joystick was used to navigate Minerva around its environment to get the appropriate data for the occupancy maps. A textured map was used due to the large area that Minerva was responsible for and to handle the large number of people expected to visit during the exhibition. Markov

²A textured map is a visual representation of the environment helping with localization.

³Odometry readings are tracking the amount of wheel turns of the robot.

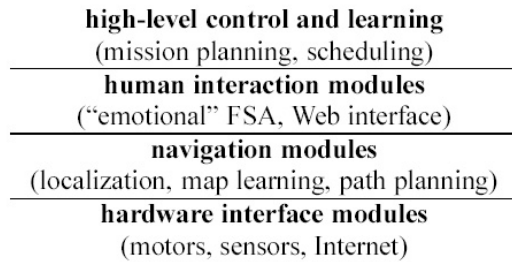


Figure 4: This is a diagram of Minerva’s software architecture.

localization was also used to filter out the genuine data readings from those corrupted by human interference.

Minerva used a collision avoidance algorithm called DWA which takes raw proximity data with the desired target location and outputs a speed and direction accordingly. A goal of this algorithm was to keep Minerva away from wide open spaces so that place land markers were available. The goal behind this was to prevent Minerva from ever getting lost.

Minerva also used a program as a way to bridge high and low level aspects of Minerva’s make up (see Figure 4). This high level control allowed Minerva to create tours on the fly which allowed Minerva to handle situations where an exhibit was overcrowded or a path was obstructed. Minerva could also monitor the execution of the tours and change the course of action for occurrences such as a low battery. Minerva also included a human interaction component so that visitors would find it approachable and entertaining. Minerva was equipped with four different emotional states which were expressed by using facial expressions and voice tone changes. Since most people had no previous experience with being around robots, it was important for Minerva to be approachable and personable.

Similar to Rhino, Minerva also had a web interface in which virtual visitors could receive tours via the Internet. Minerva was more successful than Rhino since it was able to handle



Figure 5: Swarthmore College’s Alfred.

large crowds and a larger environment. Minerva demonstrated the use of vision for more consistent navigation, the importance of obstacle avoidance, and that wide open spaces should also be avoided.

2.3 Alfred

Alfred was Swarthmore College’s entry to the American Associate of Artificial Intelligence (AAAI) national robot competition in 1999 [7]. The competition required teams to create a robot that could serve hors d’oeuvres to the conference attendees. Alfred successfully served food while recognizing humans and repeat visitors interacting with everyone in a friendly human-like manner. Alfred’s goal was to cover as much ground as possible and serve people until there was no more food on its tray and then return to the refilling station.

Due to the emphasis on covering a large amount of area, Alfred had two major modes of operation controlled by a finite state machine: move-engage and move-ignore (see Figure 6).

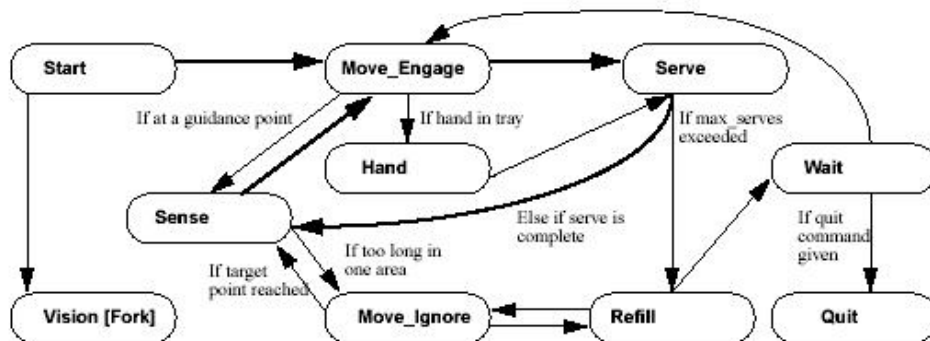


Figure 6: Description of Alfred’s finite state machine.

Move-engage was used to interact with attendees and move away from the refill station. Move-ignore was used to indicate that Alfred was either staying in one place for too long or the refilling station needed to be visited.

Vision played a large part in Alfred’s success since it relied on this to detect people, tell whether a person had previously been served, to locate the refill station, to avoid obstacles, and stay within the necessary boundaries. Blob detection ⁴ was used to recognize judges and other important people while color and texture histograms ⁵ were used to recognize repeat visitors. Alfred distinguished people from inanimate objects by using motion and skin detection with an eye template matching algorithm and color matching.

Alfred interacted with humans by using pre-recorded sound bytes and voice recognition software called ViaVoice, developed by IBM. The team was able to have Alfred understand a number of responses such as “yeah”, “sure”, and “I’ll pass” as either being affirmative or negative answers by setting up a Backus-Naur Form ⁶ file. Alfred had a more human-like quality since it used questions and answers pre-recorded by humans and could understand

⁴Blob detection looks for blocks of the same color.

⁵These are histograms measuring the frequency of a patten of color or texture.

⁶Backus-Naur is a way to express a context-free grammar used to describe a language.



Figure 7: Grace.

more colloquial responses.

Alfred's many uses of vision, especially in recognizing repeat visitors, were of great interest to us as we hope to accomplish a similar task. Finite state machines were also a technique that we thought would be necessary.

2.4 Grace

GRACE (Graduate Robot Attending ConferencE) (see Figure 7) is a robot designed to arrive at a conference, register for the event, and navigate itself to the correct location at the correct time to give a short talk about itself [9]. The team responsible for Grace was made up of members from Carnegie Mellon University, the Naval Research Laboratory, Northwestern University, Metrica Inc., and Swarthmore College. Grace was required to have state-of-the-art technologies to demonstrate a high level of intelligence and autonomy in a populated environment.

Once Grace arrived at the convention location, it was necessary to enter the building and

find the registration desk where a map of the building would be received. Since Grace did not get a map until after registering, it was necessary for Grace to know how to ask for instructions and how to interpret them. While Grace could not interpret pointing directions at the time, instructions like, “Go straight until you see the sign, then turn left,” could be parsed and executed. Grace used ViaVoice for speech recognition which converts speech into text strings. One of the challenges of reaching the registration desk was to locate the elevator and ride it. Upon reaching the correct floor, Grace had to locate the registration desk and its line. Once in the proper line, Grace had to keep up with the progressing line while respecting the idea of personal space. Grace then received a map and schedule of the conference once the desk was reached.

The map that Grace received was one that Grace had made itself from laser and odometry readings which were used in the form of an occupancy grid. Markov localization and a Markov decision process planner were used for navigation. The planner provided Grace with either a positive or negative reward based on the correctness of Grace’s actions.

Once Grace reached the proper location, Grace gave a brief talk and a PowerPoint presentation about itself. Grace used a computer animated face and text-to-speech application to convey information and expression.

Grace was able to complete the necessary tasks with minimal human assistance, integrating several interesting and difficult algorithms. Grace relied heavily on sensor readings and speech recognition to ensure that things went smoothly.

2.5 Virgil

In early 2001, a team from Rice University sought out to create a robot named Virgil that could give an outdoor tour of part of Rice's campus [10]. Virgil used GPS ⁷ and odometry readings for navigation through busy sidewalks and streets. The GPS readings were used to check the margin of error in the odometry readings since there are many factors that could cause the odometry to be miscalculated. A Kalman filter ⁸ was implemented to synchronize the odometry and GPS readings but there were times that neither of the readings were very accurate. Due to the inconsistencies of the line of sight with the GPS satellite, it was necessary for the team to use an extended Kalman filter which served as a dynamic mechanism to handle abruptions of the GPS data. GPS data was also used to calculate the error of Virgil's internal location, ensuring that the robot never veered too astray.

Most robotic projects involve the same key components: navigation, localization, and human interaction. However, each component was used and implemented very differently in each project.

3 Architecture of The Bryn Mawr Tour Guide Robot

In order to complete our goal of creating a robotic tour guide, we implemented several of the necessary components similar to those used in the projects discussed previously. We used a PeopleBot from ActivMedia which has three eight-ring sonars, tactile sensors, and a pan tilt zoom camera (see Figure 8). The robot also has built in speakers and microphones.

Our robot has three major components to its architecture which receive input from either a human user or its sensors: *navigation*, *localization*, and *speech* (see Figure 9). The navigation

⁷Global Positioning System

⁸This filter is often used to estimate the state of a system from measurements that contain random errors.




Figure 8: This is a picture of the PeopleBot robot with its indicated features.

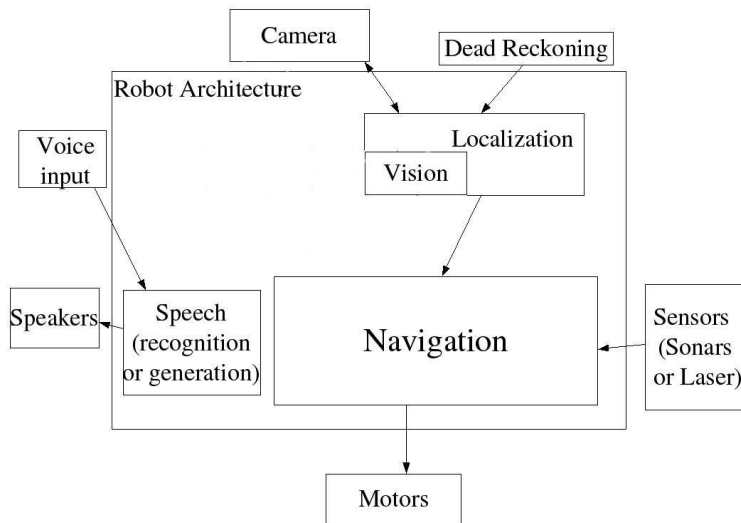


Figure 9: Representation of our robot's architecture.

module controls the robot's motors and the robot's movement. The navigation layer receives data from sensor values which provide information about the robot's current environment. Navigation also receives data from the localization component which keeps navigation in check. Localization receives data from the vision module where the robot's cameras are used to find landmarks. Once a landmark is recognized, the vision module conveys the necessary information to the localization module. The localization module synchronizes this information and helps in deciding what the robot should do next. This information is conveyed to the navigation module which then controls the motors accordingly.

The speech component receives input from a human in the form of discrete spoken words. The speech module takes the input from a human and continues to interact with the user by outputting speech through its speakers. Once the human commands are processed, they are executed accordingly.

3.1 Navigation

Navigation is a key component to any autonomous agent since it controls all robot movement. In this project, navigation uses the sonar data to determine the robot's actions as well as avoid obstacles.

The robot is able to navigate through the hallways by traveling down the middle of the hallway until a landmark is encountered. Each landmark serves as a check point for the robot to provide assurance that the robot is on the correct path.

3.1.1 Obstacle Avoidance

Our target environment is one that is frequently populated by large numbers of students and frequent obstructions in the hallways making obstacle avoidance imperative for safe

operation of the robot. Obstacle avoidance is performed by checking the sensor values to ensure that there is nothing within a certain range. On encountering/detecting an obstacle, the robot waits for a few seconds to see if the obstacle moves away on its own. If the obstacle still persists the robot turns away from it. While we intend navigation to work well enough to keep the robot away from obstacles, in general such situations need to be accounted for to prevent any accidents (i.e. bumping into walls or people).

3.2 Localization

Localization is used to give the robot an idea as to where it is in the environment as well as give direction to the robot for its next action. As mentioned previously, landmarks are used to notify the robot that it is on the correct track. We used long strips of two inch wide colored (red and blue) tape on the floor as landmarks. A vision based, color match filter⁹ is used to recognize these landmarks. The landmarks can signify a tour stop, a change of direction, or just a general check point for use in localization (see Figure 10).

3.3 Human Interaction

While the components described above are crucial for robot navigation, human interaction is needed to attract and keep an audience as well as be user friendly. Human interaction is how the robot will communicate with the robot and how humans will communicate with it. Since robotics is still somewhat of a new field for the general public and to the Bryn Mawr community, visitors and students will most likely be surprised to see this robot. The robot needs to be approachable by those who are lost as well as able to lead and entertain a group of visitors. We use speech generation and recognition such that users can interact with the

⁹A color match filter looks each pixel to see if it matches a specific color.



Figure 10: The robot in the hallway. Notice the landmarks on the floor.

robot in a ‘conversational style’.

The robot tour guide uses the speech generation software to describe and explain points of interest along the planned tour path. Speech recognition is used to get direction requests from lost users. The A* [8] algorithm is used to process the request to find the best route from the current location to the desired one. A* (A-star) is a graph search algorithm that uses heuristics to find an optimal path using straight line distances as an underestimate. The A* algorithm effectively computes the best path between any two points. We are using Festival Server for speech generation and Sphinx for speech recognition [6, 1].

4 Current Status

Over the course of this project, we have worked with two different platforms, a B21R robot with laser range finder and a PeopleBot with sonars. Since all our algorithms were imple-

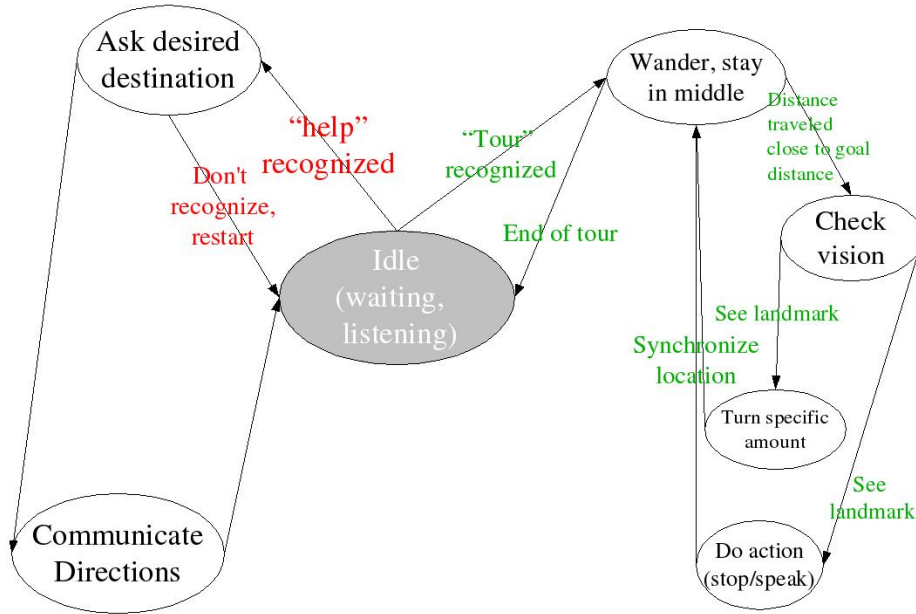


Figure 11: Diagram of our finite state machine.

mented in Pyro, minimal work was done to transition components between the two different platforms. We use a finite state machine to move from component to component (see Figure 11).

4.1 Navigation

The robot navigates itself by trying to stay in the middle of the hallways at all times. We implemented a heuristic algorithm by looking at what is ahead of the robot instead of what is immediately beside it. Obstacle avoidance is used to detect and avoid bumping into objects in the hallway. All of this was accomplished using data from the sonar sensors.

4.2 Localization

Localization is important to allow us to ensure that the robot is on track as well as providing direction to the robot if an intersection is encountered. We implement localization by using visual recognition using two different types of landmarks to signal either a tour stop or an intersection. The visual recognition module receives images from the robot's camera which is aimed at the ground.

Since vision is costly on resources, it was used only when it was needed. We use dead reckoning to determine when the robot should start using vision. Using landmarks helps correct the errors from dead reckoning and synchronizing the robot's location. Once the robot is close to an upcoming landmark, it starts using its vision to locate it. A blob filter is used to detect the landmarks.

4.3 Speech (Human Interaction)

Sphinx is an open source speech recognition application developed at Carnegie Mellon University. The Festival Client is a text-to-speech server that can convert commands and text to speech.

We limited the robot's vocabulary to include only the words or phrases that would be necessary (see Appendix A). A limited vocabulary ensures that robot will most likely recognize the words said by any user.

4.3.1 Direction Giving

The direction module consists of two key components, the A* algorithm and a converter that converts the output of the algorithm to sentences.

Since Park is a building that ironically has no logical room numbering system, navigation



Figure 12: A depiction of the points that the directs mention.

through the building is very difficult. Since one of our goals is to alleviate this confusion, we had to ensure that our directions were clear and comprehensive. With this in mind, we added a converter portion to take the output from the A* algorithm and translate the string to something that represented a more conventional form of directions.

A typical conversation with the robot may go along these lines:

Person: Help.

Robot: How can I help you?

Person: Room.

Robot: What room would you like directions to?

Person: Two Three Two.

Robot: You said room Two Three Two, is this correct?

Person: Yes.

Robot: You are at the Front Door. Go to the T Junction and turn left, go past Women's Restroom, go past Men's Restroom, go past Room 247, go past Room 233. Finally you will reach Room 232.

At this time, we are assuming that the person looking for directions will be starting at the front door. As you may notice, the phrases used by the person are very short. We specifically tried to limit the human answers to be as short as possible so that there would be less places for the voice recognition software to get confused.

4.3.2 Tour Giving

If a person requests to have a tour as the request is recognized and the robot commences the tour. Along the tour, the robot recites a script (see Appendix B) written by us from information from our Admissions office as well as various other sources. A video of the robot giving the tour is available upon request.

5 Conclusion

Although most robotics projects have been done by researchers and Ph.D. candidates, we, as undergraduates, have integrate and adapt algorithms to build a robot that gives a tour for two wings of the science building (see Figure 13). The robot recites a pre-scripted, pre-planned tour in which the robot uses dynamic navigation and localization using landmarks. Landmarks are detected using a combination of dead reckoning and vision. The robot interacts with users using speech input with a limited vocabulary. A text to speech program is used to narrate the pre-scripted tour and directions.

What was accomplished in this projects seems to have limited capabilities, however the biggest challenge is of this project is the coordination of all the components to give a tour. Obviously, each component can be improved in several ways.




Figure 13: Picture of the planned tour with the red and blue land markers in place.

One of the improvements is adapting the directions module to start from the robot's current location. With this addition, a person can stop the robot in the middle of a tour or as it is returning to its start point, the robot will be able to give directions from its current location. This will involve having the robot listen for voice interaction at all times instead of just when the robot is idle at its start point. The robot should also be able to interact with those on the tour by taking questions from the audience, so that it is more people friendly. It would be necessary to expand the range of voice recognition and provide more consistency and possibly include a natural language understanding module which will help with the robot understanding written language.

The direction module can be further improved by making the directions more concise, and yet comprehensive. At this time, the directions to a destination mention every single room from the start point to the goal. The directions can be streamlined by checking to see if there are rooms that are along the same axis such that only the rooms that have a change

in direction are mentioned. With this addition, only rooms that the user specifically needs to look for will be mentioned.

The robot should also be able to give dynamic tours in addition to the current pre-planned tour. This can be made possible by adding different lists of tour stops and offering a list that the user can choose from. For example, if the user would like to know more about the Chemistry department, the robot should be able to give a tour of the Chemistry related areas and provide more detailed information about the department. It would be necessary to expand the tour area prior to integrating this improvement which would be possible by placing more land markers and adding the correct distances and turns to the appropriate lists.

Vision could be used further by using it to recognize people as well as repeat visitors. This can be made possible by implementing texture recognition and an eye template like that of Alfred [7].

The robot could also be improved to be able to take visitors to the desired location. For example, if a person wanted to go to Room 215, the robot should be able to lead the person to the proper location or at least to the general vicinity of the goal location. The robot might also want to be able to use the elevator so that it can span the four floors accessible by the elevators.

References

- [1] A. Black and P. Taylor. Festival Speech Synthesis System: system documentation(1.1.1). <http://www.cstr.ed.ac.uk/projects/festival/manual-1.1.1/festival-1.1.1.ps.gz>.

- [2] D. Blank, L. Meeden, and D. Kumar. Python Robotics: An Environment for Exploring Robotics Beyond LEGOs. *SIGCSE*, 2002.
- [3] Douglas Blank, Holly Yanco, Deepak Kumar, and Lisa Meeden. Avoiding the Karel-the-Robot Paradox: A framework for making sophisticated robotics accessible. *AAAI Spring Symposium*, 2004.
- [4] J. Buhmann, W. Burgard, A.B. Cremers, D. Fox, T. Hofmann, F. Schneider, J. Strikos, and S. Thrun. The Mobile Robot RHINO. *AI Magazine*, 16(2):31–38, 1995.
- [5] Wolfram Burgard, Armin B. Cremers, Dieter Fox, and Dirk Hahnel and Gerhard Lake-meyer. The Interactive Museum Tour-Guide Robot. In *AAAI*, pages 11–18, 1998.
- [6] Kevin Lenzo. Sphinx Documentation. <http://fife.speech.cs.cmu.edu/sphinx/>.
- [7] Bruce A. Maxwell, Lisa A. Meeden, Nii Addo, Laura Brown, Paul Dickson, Jane Ng, Seth Olshfski, Eli Silk, and Jordan Wales. Alfred: The Robot Waiter Who Remembers You. *AI Magazine*, 1999.
- [8] Nils J. Nilsson. *Artificial Intelligence: A New Synthesis*. Morgan Kaufmann Publishers, Inc., 1998.
- [9] Reid Simmons, Dani Coldberg, Adam Goode, Michael Montemerlo, Nicholas Roy, Brennan Sellner, Chris Urmson, Alan Schultz, Myriam Abramson, William Adams, Amin Atrash, Magda Bugajska, Michael Coblenz, Matt MacMahon, Dennis Perzanowski, Ian Horswill, Robert Zubek, David Kortenkamp, Bryn Wolfe, Tod Milam, and Bruce Maxwell. GRACE: An Autonomous Robot for the AAI Robot Challenge. *AI Magazine*, 2003.

- [10] Richard Thrapp, Christian Westbrook, and Devika Subramanian. Robust Localization Algorithm for an Autonomous Campus Tour. In *ICRA*, 2001.
- [11] Sebastian Thrun, Maren Bennewitz, Wolfram Burgard, Armin B. Cremers, Frank Dellaert, Dieter Fox, Dirk Hahnel, Charles Rosenberg, Nicholas Roy, Jamieson Schulte, and Dirk Schulz. MINERVA: A Second-Generation Museum Tour-Guide Robot. In *ICRA*, 1999.

Appendix A

The words included in the robot's vocabulary include:

TOUR

ZERO

ONE

TWO

THREE

FOUR

FIVE

SIX

SEVEN

EIGHT

NINE

YES

NO

ROOM

Appendix B

START OF TOUR

Welcome to Bryn Mawr's Park Science Building. My name is Bubba and I will be your tour guide today. Please stay behind me at all times and avoid my sides if possible. Park houses classrooms and laboratories for Biology, Chemistry, Computer Science, Geology, Math, and Physics. Park also houses Collier Science Library which holds the latest science journals.

AT STAIRS/243

To the right, we have our largest lecture hall on campus. It is one of our smart classrooms and is often used for Physics lectures. To the left is a stairway that will lead you up to the Mathematics wing as well as additional Physics rooms. Math is our third most popular major with about 10% of the students majoring and is often double majored with a second discipline. Bryn Mawr has awarded degrees in math at a rate that is three times the national average for men and women.

AT PHYSICS LABS

Down this way are several Physics labs used for introductory classes. Other than MIT and CalTech, Bryn Mawr has the highest number of female Physics majors. There are additional Physics labs downstairs. Our Physics department's Professor Al Albano was selected as the most outstanding professor in Pennsylvania.

BEFORE T-JUNCTION

Bryn Mawr was the first American college to offer a Ph.D. program to women. We are currently third in the nation for the percentage of students who receive a Ph.D. in any field. About 85% of Bryn Mawr graduates go onto graduate school within five years.

AFTER T-JUNCTION

This is our Computer Science wing. Although Computer Science is not an official department and is still an independent major, it is growing in the number of current majors and professors.

PLAQUE

To the left is the plaque that three of our computer science majors won when they went to an ACM Mid-Atlantic programming competition in Fall 2003. They placed third at their specific site and thirty sixth out of 161 mid atlantic colleges and universities. Our team was one of the few all female teams to compete. Out of the seven teams from the Tri-College, our team placed second. The three students on the team were the same three people to earn a research grant to create me.

232

Down this hallway are our Computer Science labs which house roughly 50 computers. Many of the computer in the Computer Science labs run Linux, but we have a dozen which are dual boot with Windows. Computer Science recently teamed with the Physics department to create our first Beowulf cluster which is a super computer. Although our department is small, it is growing as our president Nancy Vickers has set out to expand the department. We have started the process by recently concluding a faculty search and are attracting more and more majors. About 35% of students here major in one of the sciences or mathematics and we produce a three to five ratio of the national average of female majors in the sciences. As you may have already noticed, we have a large collection of minerals. It is one of the five largest collections in the country and the collection is displayed all through the building. Our Geology department is one of the top in the country and was founded by the first woman to

hold a Ph.D. in the field.

230/231

Here is our lab where introduction to computer science students have class and do most of their programming. Here is also my home, our Emergent Intelligence lab. We have several other robots and regularly offer classes in Artificial Intelligence and Robotics in addition to a core Computer Science curriculum. We hope to one day create a robot that will be able to go from Bryn Mawr to Haverford by itself.

STAIRS/BEFORE S-JUNCTION

Up the stairs on the right is Collier Library. Combined with the other libraries on campus, Bryn Mawr has more than one million volumes. Park has been wireless since last summer so it is possible for students can work freely within the building. This is the end of my tour, thank you for joining me. I will now be returning back to the start of my tour.