

# COMP-150DR Final Project Proposal: CARoL: Coordinated, Automated Robots of Levity

**Azmina Karukappadath, Yuelin Liu, Sam Weiss**

Tufts University, COMP-150DR Developmental Robotics

**Abstract.** With the holiday season fast approaching, our group decided to focus our project around creating a caroling multi-robot system. It is an interesting topic because it not only embraces the holiday spirit, it can also be broken into sub-problems that are relevant to the course, namely: 1) door detection, 2) coordination in task completion, 3) learning rejection, and 4) developing "motivation." This project proposal will include the problem formulation, technical approach, and expected result for each of the sub-problems. We will also include an approximate schedule according to which we will work on this project.

## 1 Introduction

Developmental robotics is an interdisciplinary subject that builds upon the idea that, in order to create autonomous agents that are intelligent and adaptive, they should go through a developmental period like humans do. One key concept for development robotics is the verification principle, which emphasizes that agents cannot learn, and subsequently adaptively apply, something it cannot verify.<sup>1</sup> By making robots capable of modifying their behavior according to their experience of interacting with the environment they are situated in, they are more equipped to function in the real world, where everything in the environment is unpredictable.

For the course final project, we would like to incorporate this idea, and create a system that could learn and adaptively modify its behavior in achieving some tasks. Since the start of the holiday season coincides with the end of the semester, we thought of the idea of creating a multi-robot system that roams the second floor of Halligan and carol at each office. We wanted explore the problem of multi-robot coordination as collaboration is essential in solving real world problems of a larger scale: we could solve a problem faster if we could let a team of agents coordinate among themselves to properly divide up the problem into smaller parts, have a them each be responsible for tackling a part.

To better describe what this project envisions, we divide the problem into 4 subproblems:

1. Door Detection,
2. Coordination in Task Completion,
3. Learning Rejection, and
4. Developing Motivation

This proposal will elaborate the details concerning our project in terms of those subproblems.

## **2 Problem Formulation**

### *2.1 Door Detection*

Our project builds upon the assumption that the robots have already obtained the floor plan of the hallways of the second floor of Halligan; however, since they have to carol in front of doors, they would have to identify the locations of different office doors along the walls. Specifically, the door detection problem asks, when given a floor plan of the hallways, how a robot would learn the positions of the doors so that they know where to stop and carol.

### *2.2 Coordination in Task Completion*

Our problem formulation for the Coordination in Task Completion tasks is as follows: Given a set of doors to visit, visit all doors and carol at them in an efficient way. This means no visiting a door once it has been caroled at, as well as prioritizing doors that are most likely to have people interested in hearing carols.

### *2.3 Learning Rejection*

If there is nobody in the room, or if the person in the room does not want to be caroled at, our robots should be able to interpret "rejection," and modify its caroling behavior.

### *2.4 Developing Motivation*

In order to develop motivation in our robots, we need to find some way to quantify how successful they are at caroling. Once we have developed this metric then motivation is simply a process of maximizing that success criteria.

## **3 Technical Approach**

### *3.1 Door Detection*

Doors are one of the key elements that define an environment. Being able to recognize and interact with doors can quite literally expand the world in which an agent operates. For our project, since the robots available do not have the necessary parts for physically manipulating doors, we only concern with the robots' ability to detect doors.

There have been many different approaches to this problem in the literature. Some in turn chose an indirect approach of developing recognition capability for objects that are closely associated with doors (i.e. door knobs) to indirectly detect the presence of a door<sup>2</sup>; some employed a variation of Hough Transform to detect vertical edges in the robots' visual image<sup>3</sup>; others used ultrasound sensors to detect sudden changes in sensor readings<sup>4</sup>.

Our approach will be based on one involving Microsoft Kinect depth image camera described in Yuan et al., 2016<sup>5</sup>. Kinect depth sensing technology contains both an RGB camera and an infrared emitter and camera; therefore, it is able to capture both the color and the depth of a pixel. Nonetheless, for the purpose of detecting a door, we focus on the features provided by the infrared emitter and camera.

Using structured light, Kinect emits a known pattern of infrared light, and produces a depth image of its visual field by comparing the original pattern to the pattern distorted by the properties of the surface<sup>6</sup>. In a 2011 paper, Khoshelham demonstrated that, as the distance between the infrared camera and the surface increases, random errors of the depth measurement also increases, resulting in 0 values<sup>6</sup>. An opened door will therefore result in a depth image with a considerably large block of zero values. Further information (e.g. the corner coordinates of the door, the angle at which the door is propped open) could be extracted, but for the purpose of our experiment, simply detecting the change in depth due to an opened door would suffice.

Since the robot already has already obtained a mapping of the second floor of Halligan, we would have the robots roam around, and whenever it detects a door using depth image, simply store the coordinate information for future retrieval. Furthermore, using both the coordinate information for the position of a door and the depth sensing technology, the robots can also detect whether the door is open or closed.

### 3.2 *Coordination in Task Completion*

In order to integrate coordination and task completion with our ideas about robot motivation, we will implement a planning and coordination based around a bidding system.<sup>7</sup> In a 2006 paper, Sheng et. al. describe a system of identical and limited robots exploring an unknown space. Their solution to this problem centers around robots bidding on "rights" to explore previously unexplored areas. This proved to be an effective and distributed way for robots to partition unexplored space. It is distributed because each robot is able to determine a cost and benefit for exploring the new space. It is efficient because robots with low bids (and therefore are far away/currently busy) will lose out on auctions to closer or less occupied robots. Their paper describes an algorithm by which robots can determine their own optimal bid. We've replicated it here:

$$g_i = \omega_1 I_i - \omega_2 D_i + \omega_3 \lambda_i$$

$\omega_i$  = Some positive weights

$I_i$  = Information gain

$D_i$  = Distance to area

$\lambda_i$  = Nearness compared to other robots

This is an easy way to determine the net benefit to the robot of exploring a new area, but it comes with some serious limitations. The most serious limitation is that it doesn't account for partially completed actions. With our caroling robots, it's quite likely a robot will encounter a closed door, while another is caroling. This would initiate a new bidding process, and the above formula would fail to account for the fact that one robot is in the middle of caroling and needs to finish that action before moving to a new location. In order to account for this weakness, we propose the following

formula for computing the optimal bid for a robot caroler:

$$g_i = \omega_1 I_i - \omega_2 D_i - \omega_3 \xi + \omega_4 \lambda_i$$

$\xi$  = Seconds remaining in action queue  
 $\omega_i$  = Some positive weights  
 $I_i$  = Information gain  
 $D_i$  = Distance to area  
 $\lambda_i$  = Nearness compared to other robots

Additionally, we will redefine information gain in order to account for doors that are checked, but not caroled at. We will assume that each door that has not been checked represents an information gain of 100, each subsequent check after that will get  $\frac{1}{2}$  the information gain of the previous check. Checks after some threshold  $k$  will represent an information gain of 0.

$$I(i) = \begin{cases} 0, & \text{if } i > k \\ 100 \times (\frac{1}{2})^i, & \text{otherwise} \end{cases} \quad (1)$$

Once each caroler has computed their own bid, they will submit it to a central database where the maximal bid will be computed, again using a function proposed by Sheng et. al:

$$B = \max_i b_i$$

In order to manage this system of bidding and bid selection we will implement a centralized bidding and coordination server. This will take the form of a simple node.js based webserver or some similar technology. When caroling robots turn on and connect to the network they will connect to a previously configured IP address or URL. They will send this web server a message announcing that they are active. They will send the server another message every 5 minutes to announce that they are still active. This will allow the centralized server to manage the active robots and poll them for bids when an auction comes up.

Auctions will be initiated by robots with no current goals. A goal in our project is simply a location to check out. Each robot will maintain a queue of actions it needs to perform. This is referenced specifically in our bid computations. When a robot exhausts all of the tasks in its queue it will initiate a new auction. The auctioning server will select a door close to the free robot and then will initialize a bid on that door. If the robot that initiated the auction doesn't win that door, it will initiate a new auction until it has a goal.

Unexpected events may befall our carolers such that they are unable to complete the tasks in their queue. In order to ensure that these doors do not go uncaroled, the central planning server will also maintain a copy of each robot's queue. If a robot fails its 5 minute checkin with the server, the server will attempt to initiate communication with the robot for another 5 minutes. If these efforts fail, the robot is no longer considered part of the network, and everything in its task queue will be

auctioned off in the next auctions. Additionally, if the server asks a robot for a bid in an auction and doesn't hear back in the next 1 minute, the robot doesn't get to bid in that auction. This also triggers the 5 minutes of attempted communication and subsequent task queue abandonment.

### 3.3 Learning Rejection

Just as in traditional caroling, not every carolee appreciates the joy and merriment at that specific time. Over time, the robots will learn which professors in Halligan dislike the caroling, and eventually will stop visiting them. The robots will time the duration of time they were allowed to carol, which is marked by how long the door is kept open after knocking, to quantitatively measure how much each professor appreciates the caroling; longer open-door times indicate a greater appreciation. The robots will merge their acquired data together, so that there is a greater volume of data to learn from and so that the robots will all act in a unified manner. The likelihood that a certain professor will be visited will be represented as a percentage, which will increase and decrease as a function of how many times the robot has been rejected. We will be using Bayesian inference for this, defined as follows:

$$P(\text{successful carol} \mid \text{specific door}) = \left( \frac{\text{total successful carols at door}}{\text{total visits to door}} \right) \quad (2)$$

This probability defines the success rate of caroling at a specific door, and also defines the likelihood that the door will be visited by any of the robots. For example, a room that has been approached 100 times but only was successfully caroled into 40 times would have a probability of 0.40 that any robot will visit it in the near future. This way, if a professor is busy at the time at which they were visited by a robot, they will only be slightly less likely to be visited in the future.

In order to correctly compute this probabilistic function, the robots must have shared definitions of both successful and failed caroling attempts. We originally thought to define a successful caroling attempt as one that allowed the robot to reach the end of a song; however, we realized that the definition of successful and failed attempts will change over time according to how novel the robots are. We expect that the duration for which individuals will listen to the robots will be much higher at the beginning of the project, simply due to a sense of novelty. Over time, we expect that the duration of caroling will decrease, so we will use a learning algorithm to refine the definition of success and failure. We are currently planning to use a hyperplane algorithm that divides the time duration values of successful and failed caroling attempts.

### 3.4 Developing Motivation

We decided upon a relatively simple solution for this. In each robot, we will represent "motivation" as the probability that it will agree to take on the task to carol at a door. It will be initialized to 100%.

Whenever the robot faces rejection, whether it is because there is no one in the room to open the door, or because the person in the room does not want to be caroled at, the "motivation" probability decreases by 20%. Whenever the robot successfully caroled at a door, its "motivation" probability increases by 20%. The probability will be upper-bound at 100% and lower-bound at 10% so that it won't completely lose faith in caroling.

## **4 Expected Results**

### *4.1 Door Detection*

After the training session, we expect each robot to keep a floor plan of the second floor hallways of Halligan. It can be represented in either of the following two ways: either 1) as modifications made to the representation of the floor plan, which also encodes the representation of doors, or 2) coordinates of all valid "destinations" for caroling stored apart from the floor plan. Then, we would be able to specify any of those positions as a "destination" for a robot to go to and rotate into their caroling position facing the door.

### *4.2 Coordination in Task Completion*

The expected results of our robot coordination is three fold. The first major expected result is completion. We expect that our robots will attempt to carol at each door. Second, we expect that robots will perform in reasonable time. That is to say, robots will not drive across Halligan more than reasonable. Finally, we expect that our group of robots will be fault tolerant. Even if one of our robots disconnects from the network or fails unexpectedly every door in Halligan should be caroled at.

### *4.3 Learning Rejection*

In order to correctly learn rejection, we expect the robots to 1) develop an evolving sense of success in caroling and 2) continuously update data for professors such that professors that become disinterested in listening to carols over time are not caroled to by any of the robots. We expect that over time, they visit the doors of professors who close their doors quickly less frequently.

### *4.4 Developing Motivation*

It is very important for our carolers to be motivated and seek out caroling opportunities on their own; however, we would also like to model encouragement and discouragement observed in human beings. We expect the most rejected robot to be less likely to take on a task, while the most successful robot is the most likely to take on a task.

## 5 Schedule

**Week of October 30<sup>th</sup>:** Meet with the course staff to set up the robots.

**Week of November 6<sup>th</sup>:** Implement and finalize door detection.

**Week of November 13<sup>th</sup>:** Progress Report I Due Thursday November 16<sup>th</sup>

**Week of November 20<sup>th</sup>:** Implement and learn rejection

**Week of November 27<sup>th</sup>:** Implement multi-robot coordination

**Week of December 4<sup>th</sup>:** Implement final touches, adjust design elements

**Week of December 11<sup>th</sup>:** Finish Write-Up

## References

- 1 A. Stoytchev, "Some basic principles of developmental robotics," *IEEE Trans. on Auton. Ment. Dev.* **1**, 122–130 (2009).
- 2 S. Kim, H. Cheong, D. H. Kim, *et al.*, "Context-based object recognition for door detection," in *2011 15th International Conference on Advanced Robotics (ICAR)*, 155–160 (2011).
- 3 R. Muñoz-Salinas, E. Aguirre, and M. García-Silvente, "Detection of doors using a genetic visual fuzzy system for mobile robots," *Autonomous Robots* **21**, 123–141 (2006).
- 4 P. Carinena, C. V. Regueiro, A. Otero, *et al.*, "Landmark detection in mobile robotics using fuzzy temporal rules," *IEEE Transactions on Fuzzy Systems* **12**, 423–435 (2004).
- 5 T. H. Yuan, F. H. Hashim, W. M. D. W. Zaki, *et al.*, "An automated 3d scanning algorithm using depth cameras for door detection," in *2015 International Electronics Symposium (IES)*, 58–61 (2015).
- 6 K. Khoshelham, "Accuracy analysis of kinect depth data," in *ISPRS workshop laser scanning*, **38**(5), W12 (2011).
- 7 W. Sheng, Q. Yang, J. Tan, *et al.*, "Distributed multi-robot coordination in area exploration," *Robotics and Autonomous Systems* **54**(12), 945 – 955 (2006).