Behavior Modeling for Robot Theory of Mind

Maitry Ronakbhai Trivedi,[1] Saniya Vahedian Movahed, [1] Russell Perkins [1]
Paul Robinette [1] Seyed Reza Ahmadzadeh [1] Maria E. Cabrera [1]

[1] University of Massachusetts Lowell
MaitryRonakbhai_Trivedi@student.uml.edu, saniya_vahedianmovahed@student.uml.edu, russell_perkins@student.uml.edu,
Paul_Robinette@uml.edu, Reza_Ahmadzadeh@uml.edu, Maru_Cabrera@uml.edu

Abstract

A sizable portion of research in Human-Robot Interaction (HRI) focuses on understanding human goals from the perspective of the robot so that the robot can learn to work with people effectively. With interaction being a two-way street, it is also essential for humans to comprehend robot behavior and make informed decisions when allocating or collaborating on tasks with them. In this work, we explore the concept of the Theory of Mind (ToM) applied to decision-making in HRI. We have designed an experiment to investigate how quickly a human can recognize a robot's capabilities and how this may impact the attitudes towards or perceptions of the robot's performance.

Introduction

Humans can quickly learn to predict the actions of other people who constantly share their immediate environment. The ability to predict and understand others' behaviors helps humans to live and work together with more ease and fewer misunderstandings or errors. As robots become more pervasive in multiple fields, it becomes expected and necessary to quickly communicate the robot's attributes and capabilities to surrounding users, particularly in scenarios where there is direct human interaction and collaboration. Previous studies have shown that robots that are successfully able to predict human behavior perform better in many applications, such as assistive robotics (Losey et al. 2020),(Castelfranchi and Falcone 1998), smart vehicle resource allocation (Nicolo'Brandizzi, Brociek, and Wajda 2021; Hu et al. 2022), collaborative game (Nguyen et al. 2011), and in risk-aware decision-making which might be due to human biases toward their own or others' perspective and opinion (Kwon et al. 2020).

Even though there has been great progress toward robotic autonomous operations in multiple fields and tasks, there is still a gap where humans outperform robots in tasks where a combination between decision-making and mobility/manipulation exists. One possible reason contributing to this gap is the difference in perspectives between humans and robots, and being able to discern between possible robot behaviors and actions. Therefore, it is difficult to predict a robot's behavior while communicating with it if the conceptual model of the robot is wrong. In some cases, this misinterpretation could be very costly (for example in rescue tasks, autonomous cars, etc.) as cognitive biases toward irrational but risk-aware behaviors are prevalent in real-world scenarios. So, to ensure seamless and safe collaboration among robots and humans, robots need to anticipate how humans will behave; the opposite is equally necessary. For example, consider a scenario in which a human-driven car may encounter an autonomous car as the autonomous car makes an unprotected left turn. If the human couldn't anticipate that the autonomous car may try to make a turn even when the light might be running out, this left turn scenario may cause an accident (Kwon et al. 2020). Thus it is very important to build the best possible mental model of the robot, as it provides a formal mechanism for achieving fluent and effective teamwork during HRI by enabling awareness between teammates and allowing for coordinated action.

While mental models are organized knowledge structures that let people interact with their environment, Theory of Mind (ToM) is an ability to attribute thought, desires, and intentions to others and it is critical for day-to-day human

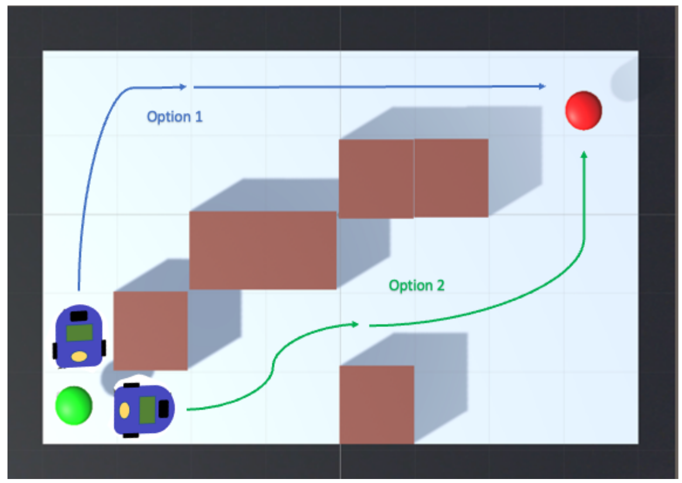

Figure 1: Top view example of 3D experiment set up. Two optional paths are shown. The path shown in option 1 (blue) is longer by one square but contains only one turn. The path depicted by option 2 (green) is shorter but contains 2 turns.

Presented at the AI-HRI Symposium at AAAI Fall Symposium Series (FSS) 2022

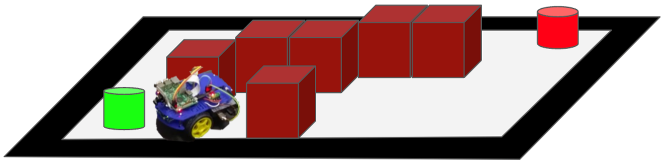

Figure 2: The user determines Duckiebot's ToM through multiple iterations in a 3D environment

social interactions such as analyzing, judging, and inferring others' behaviors, with evidence that typically developing humans exhibit this capability at age 3 (Gopnik et al. 2001). Mental modeling and ToM could lead a robot to form a belief over a human's mental model of the robot. This kind of modeling is defined as second-order mental modeling which enables robots to estimate how a human's mental model is affected by its own behavior (Brooks and Szafir 2019; Chen, Vondrick, and Lipson 2021). In general, a mental model and ToM could help in describing, explaining, and predicting the behavior of a singular system, or the behavior of a team through the idea of a shared mental model. While there has been extensive research in the understanding of human ToM (Bellas et al. 2010; Blum, Winfield, and Hafner 2018; Scassellati 2002), there has been very little research aimed toward the study of humans understanding the robot's ToM.

Through this study, we aim to convey that it is important to understand how people establish and update a robot's ability to infer the mental states of other minds (ToM) for people to make better predictions of the robot's actions and, thereby, enhance task performance in human-robot interaction settings. Therefore, we have developed an experiment with different interfaces where we investigate how quickly humans can predict the robot's capabilities and how they update their assumptions about its capability. We present a study design and discussions, as well as hypothesize how the world looks from either the robot or the human perspective.

Methodology

We have divided our experiments into two phases: 2D grid-world interface and 3D simulation. The 2D interface is created using the open-source pygame module, while the 3D simulation is created using the Unity engine. The goal of creating a 3D simulation interface is to build different experiment designs before moving on to the more challenging and time-consuming real-world experiment. Different participants will experience these two interfaces. We believe that by making the interaction environment more realistic, the participants can better understand the environmental context and the robot's intentions. Therefore, the participants will be interacting with the Duckiebot robot (Paull et al. 2017). The main goal for the participant is to predict the actions taken by the robot in a particular situation and through multiple iterations understand the ToM of the robot.

A simple user interface (UI) provides instructions about how they can play the game and what is expected from the participants. The participant can start playing the game using arrow keys and predict the robot's path during the play from the start point to the goal. Once the player predicts the path taken by the robot, they will be shown whether they succeeded at predicting the robot's action or not. Afterward, they will be shown the actual path the robot would take in the given situation.

Experiment Design

The main goal for the participants is to predict the path the robot would take in the given situation by playing the game using the keyboard. Apart from the usual up, down, left and right actions, here the participants will also be able to take diagonal actions using the Q, E, Z, and C keys on the keyboard. Therefore, the participants can choose from a total of 8 actions to reach the goal location. The participants will be playing this game for 10 rounds. We can learn from the data we collect by letting the participants play through the interface what the users believe the robot should do in a particular circumstance. The participants can use any of the 8 arrow keys to make the moves. Although they can convey 8 actions with keys, the robot's capabilities may differ from that without the participants' knowledge. The user interface is a way for the participants to express their thought on what path the robot would take but it is not necessary that robot would follow the same path due to different capability. The participants will not be informed about what kind of actions the robot will be able to take prior to the experiment. After each round, the correct demonstration of the robot's actions will be shown to the participants. This will help them analyze what are the possible actions that can be taken by the robot and make adjustments to their decisions in the next round.

We expect the participants to predict what actions the robot can take and understand their capability by watching this demonstration. In every round, the initial and the goal location will be changed to avoid redundancy. After the completion of the game, the participants will be asked to fill out a questionnaire about whether they understood the robot's actions and if they were able to comprehend the robot's capabilities.

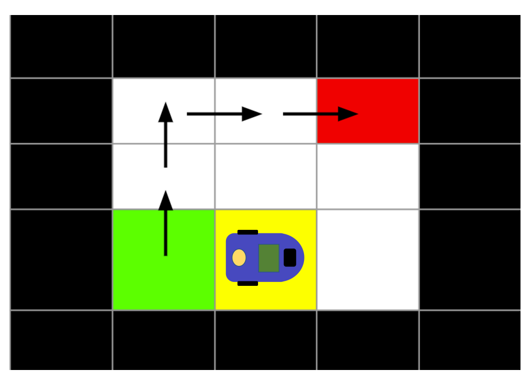

Figure 3: The correct trajectory performed by the robot

2D grid-world interface

As shown in Figure 3, the online participants for the first part will be interacting with a 5x5 grid followed by a 7x7 grid, where the green colored box is the initial position of the robot and red colored box is the goal position of the robot. The robot's current position will be displayed as a yellow box. The black colored boxes are the spots where the robot will not go or is not allowed to go. The main goal for the player is to take the robot to the goal location with minimum steps. Here, the number of cells visited by the robot will be considered as the number of steps taken by the robot.

In the 2D scenario, the participants can choose from the 8 actions while the robot does not have the ability to take the diagonal action. Therefore, an example showing the correct path the robot will be taking in the given situation is shown in Figure 3. Initially, the participants will be lacking the ToM for the robot and they will be assuming that it is obvious that the robot can also make diagonal movements and it will only take two steps to reach the goal. After playing some number of rounds and watching the path taken by the robot, we expect that the participants will realize that the robot is not able to take the diagonal actions and it will take 4 steps to reach the goal location.

3D simulated environment

For the 3D scenario, we have developed a 3D simulation game with obstacles using Unity. There is a fixed number of obstacles, each obstacle takes up one grid square. The example of the 3D environment is shown in Figures **??** and 4. The 3D simulated environment has the same instruction block in the UI as the 2D environment. The grid size is 5x8 instead of 5x5. A top-down view of the 3D environment is shown to the user at the same time, as shown in Figure **??**. The map will not contain the path routes. When the game begins the user will see the 3D environment.

The UI will facilitate the movement of a Duckiebot robot conducting navigation tasks. In each round, there will be a simulation of the robot's capabilities. After the simulation is complete, the interface will show an instruction block with the actions available to the user. The UI will indicate that it is time for the user to start moving the robot and a timed countdown is displayed. At the end of the countdown, the UI will compare the user's actions and the robot's actions by attributing a score based on the amount of distance traveled and the number of turns used. The size of the environment could increase so that the number of path options for the robot increase and there can be a "live" run that would simulate the user controlling the robot in real-time. In this scenario, a mistake would result in collisions with the environment.

After the participant completes each round of the 3D simulation, they will be observing a Duckiebot navigate a real course. The path that takes the least time is the one with the least turns. In this scenario, the important factor to observe is the total time taken by the participants and the robot to complete the task. The Duckiebot does not turn very efficiently, the reasons for this are covered in the Robot Overview section of the paper. Due to the difficulty in turning, we can say that when the participants successfully navigate the environment by creating a path that has the least amount of turns, one, they have understood the robot's theory of mind.

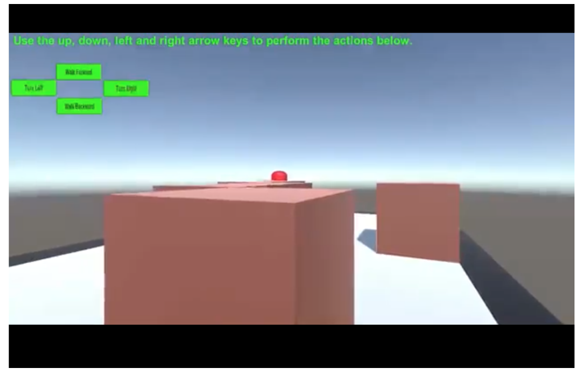

Figure 4: First person view from Duckiebot's camera in 3D environment

Hypothesis

In the 2D phase of the experiment, the robot's mental model has four actions and there are eight action options available to the participant's mental model through the UI.

> **H1:** The user's ability to predict the robot's action will decrease proportionally to the difference in the robot's capabilities such as motion.

Moreover, if the complexity of the environment increases (i.e. from a 5x5 to a 7x7 grid) while keeping the action options available to the user and the robot the same, it will take longer for the humans to predict the robot's behavior.

> **H2:** A person's ability to predict the robot's action will be inversely proportional to the complexity of the environment.

Regarding the environmental context, we submit that by having an environment that more closely resembles the real world, participants will be able to predict the robot's ToM faster than in the 2D scenario.

> **H3:** A user in a 3D simulated environment and watching a real robot will be able to understand the robot's ToM faster than in the 2D experiments.

Metrics

In order to test our hypotheses, we have considered the following metrics and parameters:

- Number of Squares: We track the number of steps the participants used in order to reach the goal. Each step corresponds to the square in the grid excluding the initial location cell. The optimal number of steps the robot takes is 4. Therefore, we can say that when the user successfully completes the task in 4 steps, it understands the robot's ToM.
- Number of Diagonal moves: The participants can make diagonal moves using the Q, E, Z, and C keys. They are expected to realize that the robot is not capable of taking diagonal moves once they watch the video of the robot performing the task. Therefore, if they take even one diagonal move, it is not acceptable.
- Time: We assume that it takes 0.1 seconds for the participants to take one step. Therefore, we can say that when

the participant completes the game in 0.3 seconds, they have understood the robot's ToM.

- Survey Questions: We have prepared a small questionnaire for the participants regarding their experience playing the game. we use a scaled score method (a number within 1-5).

Moreover, we observe the number of rounds it takes for the participants to understand the robot's capabilities. We assume that for a small environment of a 5x5 grid, it will not take more than 4 rounds to do this. For the 5x8 grid, it will take no more than 8 rounds. As the complexity of the environment increase. it will take longer for the participants to predict the robot ToM which proves our hypothesis 2.

Robot Overview: Duckiebot

For this experiment, we chose to use the Duckiebot robot made by Duckietown. It is a differential drive robot with the following in the kinematic model.

$$\frac{dx}{dt} = \frac{V_R + V_L}{2} cos\omega \tag{1}$$

$$\frac{dy}{dt} = \frac{V_R + V_L}{2} sin\omega \tag{2}$$

$$\frac{\omega}{dt} = \frac{V_R - V_L}{L} \tag{3}$$

The Duckiebot has Hall effect sensor wheel encoders, a front-facing 160° FOV camera, an internal measurement unit (IMU), and a front-facing time of flight sensor. It also has a 2G NVIDIA Jetson nano for computation.

For this experiment, we assume that the camera has been calibrated so that there are no issues with obstacle identification and that the odometer has been calibrated so that there is no drift while driving in a straight line.

The Duckiebot can be manually controlled for with a keyboard controller. The arrow keys are used to steer the robot. The gain for the velocity of the robot is fixed. The forward arrow on the keyboard will move the Duckiebot forward. Since the steering is calibrated, it is assumed that the robot will travel in the direction of its heading with no error. The right and left arrows turn the robot in the respective directions. The steering is sensitive because the velocity is fixed, and the robot will often overshoot the desired heading when the keypad is used to steer it. The best technique to achieve an accurate heading is to use small presses of the correct arrow key until the robot achieves the correct heading. This of course takes time.

The amount of time it takes to turn is much larger than the amount of time that it takes to travel straight. Therefore, the Duckiebot's mental model is represented by the choice of a path that minimizes turns rather than choosing the shortest path. The algorithm used is for path planning is presented in (Hassani, Maalej, and Rekik 2018).

Discussion

Given that the hypothesis' are confirmed, these experiments would demonstrate a person's capacity to comprehend a robot's mental model through observation and trial and error. In the first experiment, the robot's capability was limited to moving straight and it could not move diagonally. Humans had the ability to move diagonally and would naturally default to the diagonal movement. This is because of the fact that over time humans learn to find the shortest path and they might think that the robot also has the capability to move diagonally. as they do. After some rounds of observations, the human would realize that that capability did not exist in the robot.

In the 3D experiment, the human had the ability to move the robot with the arrow keys on the keyboard. Once again human is expected to default to the shortest path. In this case, Duckiebot's mental model was a path-planning algorithm that minimized the number of turns. The human could understand kinematics described in Equations 1, 2, and 3 but hidden from them was the fact that the time it took to achieve the correct heading when there was $\frac{\omega}{dt}$ was much greater than when the Duckiebot's velocity was straight $\frac{dx}{dt} + \frac{dv}{dt}$.

Due to the increased complexity of the experiment as well as the ability to freely use the arrow keys to execute turns and forward movement it would take the human more time to understand the mental model of the robot. This is important because as intelligent autonomous robots are designed the relative complexity of the task can be used as a metric for how much information needs to be supplied to a human operator or teammate so that optimal performance can be achieved.

One of the limitations of this study is that hypothesis H3 assumes that the introduction of 3D obstacles will have no effect on the path selection. In the future, we would like to extend the experiments where the participants will be collaborating with an advanced robot performing a more complex task.

References

Bellas, F.; Duro, R. J.; Faiña, A.; and Souto, D. 2010. Multilevel darwinist brain (mdb): Artificial evolution in a cognitive architecture for real robots. *IEEE Transactions on autonomous mental development*, 2(4): 340–354.

Blum, C.; Winfield, A. F.; and Hafner, V. V. 2018. Simulation-based internal models for safer robots. *Frontiers in Robotics and AI*, 4: 74.

Brooks, C.; and Szafir, D. 2019. Building second-order mental models for human-robot interaction. *arXiv preprint arXiv:1909.06508*.

Castelfranchi, C.; and Falcone, R. 1998. Towards a theory of delegation for agent-based systems. *Robotics and Autonomous systems*, 24(3-4): 141–157.

Chen, B.; Vondrick, C.; and Lipson, H. 2021. Visual behavior modelling for robotic theory of mind. *Scientific Reports*, 11(1): 1–14.

Gopnik, A.; Sobel, D. M.; Schulz, L. E.; and Glymour, C. 2001. Causal learning mechanisms in very young children: two-, three-, and four-year-olds infer causal relations from patterns of variation and covariation. *Developmental psychology*, 37(5): 620.

Hassani, I.; Maalej, I.; and Rekik, C. 2018. Robot Path Planning with Avoiding Obstacles in Known Environment Using Free Segments and Turning Points Algorithm. *Mathematical Problems in Engineering*, 2018.

Hu, T.; Ma, H.; Liu, H.; Sun, H.; and Liu, K. 2022. Self-attention-based Machine Theory of Mind for Electric Vehicle Charging Demand Forecast. *IEEE Transactions on Industrial Informatics*.

Kwon, M.; Biyik, E.; Talati, A.; Bhasin, K.; Losey, D. P.; and Sadigh, D. 2020. When humans aren't optimal: Robots that collaborate with risk-aware humans. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 43–52. IEEE.

Losey, D. P.; Srinivasan, K.; Mandlekar, A.; Garg, A.; and Sadigh, D. 2020. Controlling assistive robots with learned latent actions. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 378–384. IEEE.

Nguyen, T.-H. D.; Hsu, D.; Lee, W.-S.; Leong, T.-Y.; Kaelbling, L. P.; Lozano-Perez, T.; and Grant, A. H. 2011. Capir: Collaborative action planning with intention recognition. In *Seventh Artificial Intelligence and Interactive Digital Entertainment Conference*.

Nicolo'Brandizzi, S. R.; Brociek, R.; and Wajda, A. 2021. First Studies to Apply the Theory of Mind Theory to Green and Smart Mobility by Using Gaussian Area Clustering.

Paull, L.; Tani, J.; Ahn, H.; Alonso-Mora, J.; Carlone, L.; Cap, M.; Chen, Y. F.; Choi, C.; Dusek, J.; Fang, Y.; Hoehener, D.; Liu, S.-Y.; Novitzky, M.; Okuyama, I. F.; Pazis, J.; Rosman, G.; Varricchio, V.; Wang, H.-C.; Yershov, D.; Zhao, H.; Benjamin, M.; Carr, C.; Zuber, M.; Karaman, S.; Frazzoli, E.; Del Vecchio, D.; Rus, D.; How, J.; Leonard, J.; and Censi, A. 2017. Duckietown: An open, inexpensive and flexible platform for autonomy education and research. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*, 1497–1504.

Scassellati, B. 2002. Theory of mind for a humanoid robot. *Autonomous Robots*, 12(1): 13–24.