Modeling Human Adaptation in Repeated Collaborative Tasks

Stefanos Nikolaidis
Carnegie Mellon University
snikolai@cmu.edu

Siddhartha Srinivasa
Carnegie Mellon University
siddh@cmu.com

ABSTRACT

This short paper summarizes a game-theoretic model of human partial adaptation to the robot. We model the human as following a best-response strategy to the robot action, based on theor own, possibly distorted, reward function. The model allows the robot to take *informative* actions, in order to teach the human its capabilities.

CCS CONCEPTS

• **Human-centered computing** → *User studies*; • **Theory of computation** → *Algorithmic game theory*; • **Computing methodologies** → *Robotic planning*;

KEYWORDS

Algorithmic game theory, robotic planning, user studies

1 INTRODUCTION

A lot of work in robotics has focused on making robots part of our everyday life, helping people as effective members of human-robot teams. This requires human teammates to know what the robot can and cannot do: the robot's perceived capability should match its true capability. Otherwise, the gap in expectation can significantly reduce human-robot team performance [1].

2 PARTIAL ADAPTATION

For example, we consider the table-clearing task illustrated in Fig. 1. The user and the robot are tasked with clearing the table by placing items in the bins. The clearing task is repeated a number of times. The user lacks the following information about the robot:

- The robot does not know where the green bin is.
- The robot cannot lift the bottle that is farthest away from its base.

Using the table-clearing task described above as an example, we let the robot attempt to grasp the bottle that is closest to its base, dropping the blue bin in the process. This will likely cause the human teammate to change her actions: in the next round, she will move the green or blue bin out of the robot's way. However, without observing the robot fail in lifting the other bottle, she still has no information about which action to take (i.e. emptying the bottle of water), if the robot attempts to lift the bottle.

This is an example, where the human may change her actions based on the robot actions, while not completely adopting the robot's optimal policy. We call this *partial human adaptation to the robot.*

Figure 1: Top: User performs a repeated table-clearing task with the robot. The robot fails intentionally in the beginning of the task, in order to reveal its capabilities to the human teammate. Bottom-left: The robot drops the blue bin off the table while moving towards the left bottle. Bottom-right: The torques applied exceed their limits when the robot attempts a grasp at an extended configuration, and the robot stops moving.

3 APPROACH

In [2], we proposed a game-theoretic model of human partial adaptation to the robot. We modeled the human as following a best-response strategy to the robot action, based on their own, possibly distorted, reward function. The human reward function changed over time, as the human observed the outcomes of the robot and her own actions.

The model allowed the robot to *reason over how the human expectations of the robot capabilities would change based on its own actions*. The robot used this model to compute an optimal policy, which enabled it to decide optimally between *revealing information* to the human and *choosing the best action given the information that the human currently has.*

PETRA '17, Island of Rhodes, Greece

978-1-4503-5227-7/17/06…$15.00
DOI: http://dx.doi.org/10.1145/3056540.3076211

4 RESULTS

In [2], we proved that under certain observability assumptions, the optimal policy can be computed in time linear to the number of robot actions and the time horizon. Additionally, we showed through a human subject experiment that the proposed model significantly improved human-robot team performance, compared to policies that assume complete human adaptation to the robot. Finally, we showed through simulations that the proposed model performed well for a variety of randomly generated tasks. This work was the first step towards modeling the change of human expectations of the robot capabilities through interaction, and integrating the model into robot decision making in a principled way.

REFERENCES

[1] Victoria Groom and Clifford Nass. 2007. Can robots be teammates?: Benchmarks in human–robot teams. *Interaction Studies* 8, 3 (2007), 483–500.

[2] Stefanos Nikolaidis, Swaprava Nath, Ariel D. Procaccia, and Siddhartha Srinivasa. 2017. Game-Theoretic Modeling of Human Adaptation in Human-Robot Collaboration. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction (HRI '17)*. ACM, New York, NY, USA, 323–331. DOI: https://doi.org/10.1145/2909824.3020253