Learning to Assess the Cognitive Capacity of Human Partners

S. M. al Mahi, Matthew Atkins and Christopher Crick

Oklahoma State University, Stillwater OK 74074

smahi@okstate.edu, matthew.atkins@okstate.edu, chriscrick@cs.okstate.edu

ABSTRACT

We demonstrate a robotic system that learns to recognize the behavioral indicators that a complex, rapidly-evolving task has exceeded the cognitive capacity of a human partner. Based on that determination, it can act autonomously to reduce the human decision burden, significantly improving task performance.

1. INTRODUCTION

One of the most challenging obstacles facing human-robot teams is the inherent communication barrier between the two. Human operators, at least once they have received training, have some notion concerning the capacities of their mechanized partners, but the ability of robots to assess the limitations of humans has not received adequate attention.

In our system, the robot learns to model the relationship between human direction and task performance for a well-understood task—in this case, navigating a maze. The robot then participates in a different, more difficult problem, but it can still use its learned model to evaluate a human operator's cognitive load. A robot's ability to participate constructively in a human-robot team will benefit immensely from understanding and accommodating this cognitive stress appropriately [1]. Our work demonstrates robots that can detect the emergence of cognitive stress in their operators, increasing their level of autonomy and reducing demands on the operator's attention.

Human-robot interactions can be evaluated using fundamental metrics [5]. We leverage this data to inform our robots' estimation of a human operator's cognitive capacity. Recent work [3, 4] presented a model for assessing a human's attention level, based on eye contact and gaze detection towards a robot. In our work, the robot learns a general behavior model to identify the operator's cognitive threshold, rather than relying on the specifics of gaze.

2. PROBLEM STATEMENT

Take $H = [h_1, h_2, \dots, h_m]$ to be a vector of ecologically

valid measurements of human behavior relevant to the problem space. Assume a task for which a robot participant can independently calculate s , a function of a vector of measurable environmental features $E = [e_1, e_2, \dots, e_n]$. Thus, $s = f(E)$, where f is a task-specific function known to the robot. Using f and calculating s , a robot can build its own supervised training set for a learning task, where the human input H is associated with s through a learned function g . Thus, the robot learns to associate the human behavioral metrics H with task success s within a known task, so the output of g is a learned *estimate* of the true success ($\hat{s} = g(H)$). Now, assign the robot a task which requires human input for success, i.e., the robot has no access to an analogue to f or s in this new task. However, it can still measure the components of H , and it has access to its learned model g . We show that computing $\hat{s} = g(H)$ in this new environment allows the robot to estimate not the task success, but the cognitive load on its human partner and an estimate of the quality of the human's direction.

3. EXPERIMENTAL DESIGN

Our experiments consisted of two games, maze navigation [2] and coin collection. All the games were played in two configurations, using either one or two robots, and with an interaction duration of two minutes. In the maze game, the robot collects data needed to build a model g for evaluating human cognitive load based on input H . In the subsequent coin game, the robot is placed in a different scenario where it has no access to success measures or even rules. Even so, with no independent means of measuring task success, it can still calculate $\hat{s} = g(H)$, and can therefore evaluate the quality of instruction, and hence the cognitive capacity, of its human partner.

In the maze game, the vector of environmental measurements E consists of the following components: e_0 is the *disparity* term, the distance between the navigation directions provided by a human and the route that the robot would have planned for itself, e_1 is the *collision* term, which penalizes collisions with walls, and e_2 is the *time delay* term, the amount of time taken for the human to guide the robot through the maze, compared with the robot's estimate of the time it would have taken under its own power. The computation of $s = f(E)$, the function for measuring success of the human directions, is a normalized summation of the elements of E .

By computing this value s , the robot can label its own data in order to train a supervised learning algorithm which will relate the success of a human-directed task with a set

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

HRI '17 March 6–9, 2017, Vienna, Austria

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4885-0/17/03.

DOI: <http://dx.doi.org/10.1145/3029798.3038430>

Figure 1: Robot’s estimated cognitive stress level modestly correlates with physiological metrics.

of measured behaviors H : h_0 is the *decision interval* term, which measures the time elapsed between the robot reaching a navigation goal and the human providing a new one, h_1 is the *error correction* term, which measures the tendency of a human operator to provide a navigation goal and then subsequently provide another before the task is complete, and h_2 is the *franticness* term, which characterizes erratic behavior for the control inputs. The robot’s model incorporates the data learned from all participants.

4. RESULTS

In general, the robot correctly predicts an operator’s cognitive load. Figure 1 shows modest correlation between physiological evidence (breathing rate measured with a Bioharness) of an operator and the robot’s estimation of stress. This is suggestive but not conclusive; it may be that physiological stress measures are not precisely indicative of the cognitive load which our robots attempt to predict.

Much more convincing is the learned model’s contribution to task success. The coin game requires the operator to navigate the maze collecting coins (visible to the human operator but not to the robot). Delays and errors in successfully collecting coins increase an operator’s task penalty score; as time pressure and the number of robots participating in the game grows, the operator’s cognitive load is likewise expected to increase. In the manual test condition, the robots continue to act according to human instruction regardless of their model’s estimate of cognitive load, while in the autonomous assistance mode, the robots revert to maze navigation behaviors whenever their learned human behavior models detect high cognitive stress. As shown in Figure 2, this behavior significantly enhances the overall performance in the game. Robots are able to reliably assess the cognitive load their human partners are under, even in contexts where the actual tasks they are being asked to perform are opaque to the robot.

5. CONCLUSION

Robots that are capable of understanding the cognitive

Figure 2: Coin game task penalties in manual vs. autonomous assistance modes across 34 test subjects. $p < 0.05$ in both instances.

load their operator is experiencing are vital to safe and efficient teamwork in complex scenarios where the proper level of autonomy and interaction is fluid. Vital communication cues are embedded in the way we behave in particular circumstances, and these implicit indicators do not have to be lost on our robots. Our work’s contribution is to demonstrate a quantitative, learnable, generalizable model that allows a robot to determine that a user has succumbed to cognitive stress, even when it cannot independently assess the instructions it is being given.

This work was supported by NSF award #1539070 (Unmanned Aircraft System for Atmospheric Physics).

6. REFERENCES

- [1] J. W. Crandall, M. Goodrich, D. R. Olsen Jr, C. W. Nielsen, et al. Validating human-robot interaction schemes in multitasking environments. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 35(4):438–449, 2005.
- [2] C. Crick, S. Osentoski, G. Jay, and O. C. Jenkins. Human and robot perception in large-scale learning from demonstration. In *Proceedings of the 6th international Conference on Human-Robot Interaction*, pages 339–346. ACM, 2011.
- [3] D. Das, M. M. Hoque, Y. Kobayashi, and Y. Kuno. Attention control system considering the target person’s attention level. In *Proceedings of the 8th ACM/IEEE International Conference on Human-robot Interaction*, pages 111–112. IEEE Press, 2013.
- [4] M. M. Hoque, T. Onuki, D. Das, Y. Kobayashi, and Y. Kuno. Attracting and controlling human attention through robot’s behaviors suited to the situation. In *Proceedings of the Seventh Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI ’12*, pages 149–150, New York, NY, USA, 2012. ACM.
- [5] D. R. Olsen and M. A. Goodrich. Metrics for evaluating human-robot interactions. In *Proceedings of PERMIS*, volume 2003, page 4, 2003.