Modeling Human-Robot Interactions as Systems of Distributed Cognition

Chien-Ming Huang and Bilge Mutlu
{cmhuang, bilge}@cs.wisc.edu
Department of Computer Sciences
University of Wisconsin–Madison

Abstract

Robots that are integrated into day-to-day settings as assistants, collaborators, and companions will engage in dynamic, physically-situated social interactions with their users. Enabling such interactions will require appropriate models and representations for interaction. In this paper, we argue that the dynamic, physically-situated interactions between humans and robots can be characterized as a system of *distributed cognition*, that this system can be represented using probabilistic graphical models (PGMs), and that the parameters of these models can be learned from human interactions. We illustrate the application of this perspective in our ongoing research on modeling dyadic referential communication.

Introduction

Following the traditional view of human cognition primarily as *internal* processes and representations, early research in artificial intelligence and robotics sought to achieve computational intelligence by developing similar internal representations and processes for planning and decision-making. While this development enabled robots to plan and perform highly complex tasks in controlled environments, the integration of robots into the human environment introduces unique challenges that stem from the dynamic, physically-, socially-, and culturally-situated interactions in which they will be engaged.

In an effort to address the limitations of the traditional view of human cognition in modeling dynamic, situated cognitive processes, the theory of *distributed cognition* observes that cognitive processes "in the wild" are usually distributed across members of a social group, involve interactions with the environment, and are distributed through time (Hutchins 1995). In this paper, we characterize human-robot interaction as a system of distributed cognition and argue that this characterization provides an appropriate basis for the use of artificial-intelligence techniques to model human interactions and to control robot behaviors in order to achieve fluent and effective interactions between humans and robots. Following a brief introduction of the theory of distributed cognition, we illustrate how human-robot interactions might be characterized as systems of distributed cognition and how this characterization serves as a basis for modeling interactions using probabilistic graphical models (PGMs) using our ongoing work on modeling dyadic referential communication.

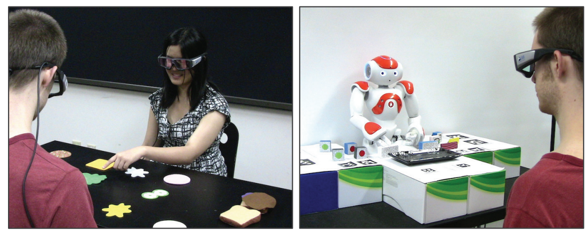

Figure 1: In our ongoing work, we characterize human referential communication (left) as a system of distributed cognition the parameters of which we learn using artificial-intelligence techniques. The trained model is used to predict human behaviors and to control robot behaviors in human-robot referential communication (right). In the task, one participant (server) prepares sandwiches according to the instructions of the other participant (customer).

Human-Robot Interactions as Systems of Distributed Cognition

The theory of distributed cognition describes human cognition as being distributed across three key dimensions (Hollan, Hutchins, and Kirsh 2000):

1. The members of a social group,
2. Internal and external (material or environmental) structure,
3. Time periods such that earlier events shape later events.

These observations extend the traditional view of human cognition primarily as *internal* processes and representations, providing a basis for characterizing dynamic, physically-, socially-, and culturally-situated human interactions. Previous research in human-computer interaction has drawn on such characterizations as a basis for exploring new forms of interaction and designing novel computer technologies (Hollan, Hutchins, and Kirsh 2000).

Distributed cognition can also be used to characterize the rich space of interactions between humans and robots. An individual's interactions with a robot and the physical environment over a period of time can be considered as instances

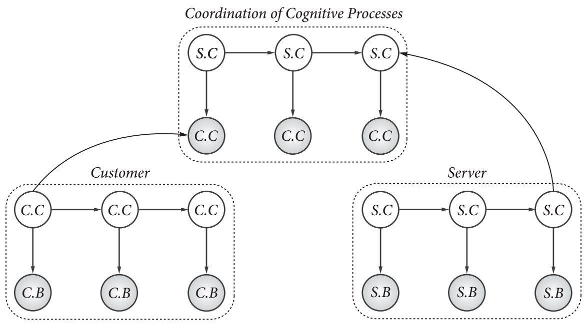

Figure 2: We construct three hidden Markov models (HMMs) to model the distribution of cognitive processes in the sandwich-making task. The bottom HMMs represent the relationship between cognitive states and behavioral cues for the server and the customer. The top HMM captures the patterns in which the cognitive processes of the two parties interrelate.

of distributed cognition. This characterization in turn can guide the design of robot systems that achieve more fluent interactions with their users by integrating their users' perspectives, knowledge, and perceptions; physical and cultural contextual factors; and the history of their interactions with their users and the environment into their planning, decision-making, and behaviors.

An Illustrative Example: Dyadic Referential Communication

Consider a scenario in which a robot takes orders from and prepares food for a customer. Fluent interaction between the robot and its user requires the distribution of a number of cognitive processes, including exchanging of information, making references to objects in the environment, and turn-taking in communication and performing the task. To model the cognitive processes involved in this scenario, we recruited 13 pairs of participants to perform in a similar food-serving task. The task involved one participant making a sandwich based on instructions from the other participant, as shown in Figure 1. Participants were instructed to add a total of 15 arbitrary toppings to the sandwich.

Participants performed the task twice, switching roles, which yielded 26 episodes of interaction. The episodes were videotaped, and participants' gaze behaviors were captured using a dual mobile eye-tracker setup. The data were annotated for high-level cognitive states, such as "thinking which topping to order," and for low-level behavioral cues, such as gaze target and speech referent, in order to model the relationship between cognitive states and behavioral cues and the patterns in which the cognitive processes of the two participants interrelated.

Building on prior work (Huang and Mutlu 2014; Lee et al. 2013; Mead and Mataric 2014), we use PGMs, such as hidden Markov models (HMMs) and dynamic Bayesian networks (DBNs), to construct representations of the relationships described above.

To model dyadic referential communication in the sandwich-making task, we used HMMs to represent the system of distributed cognition and to learn the parameters of this system (Figure 2). Two separate HMMs were used to model the relationships between the high-level cognitive states and the low-level behavioral cues for the two interaction partners. An additional HMM was used to model how the cognitive processes of the two partners were interrelated.

Once model parameters are learned from the annotated data of human interactions, they can be used to control the behaviors of a robot during its interactions with its user. The robot can infer its user's cognitive states based on its observations of its user's behavioral cues. The inferred user states can then be used to determine what state the robot should be in and what behaviors it should display in response to its user's state.

Conclusion

The increasing integration of robots into the human environment introduces new computational and design challenges in representation, recognition, and control in dynamic, physically- and socially-situated interactions. In this paper, we argued that the integration of the theory of distributed cognition from cognitive sciences and representation and learning techniques from artificial intelligence hold promise in addressing these challenges. To illustrate this integration, we presented our ongoing work that uses HMMs to model human referential communication and to control robot behaviors to achieve fluent human-robot interactions.

Acknowledgments

We thank Sean Andrist and Allison Sauppé for their help with this work. This work was supported by National Science Foundation awards 1017952 and 1149970 and the National Institutes of Health award 1R01HD071089-01A1.

References

Hollan, J.; Hutchins, E.; and Kirsh, D. 2000. Distributed cognition: toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction (TOCHI)* 7(2):174–196.

Huang, C.-M., and Mutlu, B. 2014. Learning-based modeling of multimodal behaviors for humanlike robots. In *Proceedings of the 9th annual ACM/IEEE international conference on Human-Robot Interaction*, 57–64. ACM.

Hutchins, E. 1995. How a cockpit remembers its speeds. *Cognitive science* 19(3):265–288.

Lee, J. J.; Knox, W. B.; Wormwood, J. B.; Breazeal, C.; and DeSteno, D. 2013. Computationally modeling interpersonal trust. *Frontiers in psychology* 4.

Mead, R., and Mataric, M. J. 2014. Probabilistic models of proxemics for spatially situated communication in hri.