Editorial: Representation Learning in HRI

The Editors-in-Chief and Special Issue Guest Editors welcome you to this concluding issue of Volume 8 for ACM THRI. This issue marks our second full volume with ACM Publications as its current flagship for robotics and HRI research. Our editorial team looks forward to many superb volumes to come with the rapid growth and advancement of our field. In this issue, we are pleased to feature a survey article and a selection of four papers from the nascent area of Representation Learning and its applications to Human and Robot Cognition.

Our lead-off article by Venture and Kulic, “[Robot Expressive Motions: A Survey of Generation and Evaluation Methods](#),” presents a survey on generating expressive whole body movements in physical robots and avatars. Existing ideas for embodied motion across wheeled, legged, and flying systems are first described and then synthesized into metrics that can be used to evaluate generated movements. This article further defines current open challenges for creating more acceptable and expressive motions in artificial agents, which present potential future directions for research and development in HRI.

REPRESENTATION LEARNING FOR HUMAN AND ROBOT COGNITION

Intelligent robots are rapidly moving to the center of human environments. They collaborate with human users in different applications that require high-level cognitive functions. Such cognitive abilities allow them to understand and learn from human behavior within different HRI contexts. To this end, a persistent challenge of considerable importance is that of representation learning. Emerging from artificial intelligence research, representation learning refers to learning representations of data to efficiently extract relevant features for classification, regardless of whether such classifiers are probabilistic, discriminative, or connectionist. This active area of research spans different fields and applications, including speech recognition, object recognition, emotion recognition, natural language processing, language emergence, and development.

In addition, representation learning allows for emulation of different human cognitive processes through appropriate computational modeling. Learning, in this construction, constitutes a basic operation in the human cognitive system and developmental process, where perceptual information enhances the ability of the sensory system to respond to external stimuli through interaction with the environment. This learning process depends on the optimality of features (representations of data) that allow humans to make sense of everything they feel, hear, touch, and see in the environment. Using intelligent robots could shed light on the underlying mechanisms of representation learning and its associated cognitive processes. Consequently, such understanding can bring our field closer towards making robots able to better collaborate with human users.

This special section on Representation Learning for Human and Robot Cognition aims to shed light on the cutting-edge lines of interdisciplinary research in artificial intelligence, cognitive robotics, and human-robot interaction focusing on representation learning. With a focus on creating natural and intelligent interaction between humans and robots, this special section conducted a thorough and detailed review process that produced the following four articles:

In “[The Impact of Different Levels of Autonomy and Training on Operator’s Drone Control Strategies](#),” Zhou et al. present an experimental study examining the effects of varying the levels

© 2019 Copyright held by the owner/author(s).

2573-9522/2019/12-ART19

<https://doi.org/10.1145/3366621>

of autonomy for Unmanned Aerial Vehicles (UAVs) in the process of training operators. This study used a Hidden Markov Model representation to examine two different drone control modes for operators with varying levels of supervisory control and enhanced autonomy. The results show that operators with both supervisory and enhanced teleoperation control have significant differences in their performance from operators trained in the supervisory control mode only.

Gärdenfors, in “[Using Event Representations to Generate Robot Semantics](#),” posits a geometric approach to representing concepts from human-robot dialogic interaction, such as with an iCub Robot. This approach uses events, modeled as vectors in metric “concept spaces,” to represent actions and their effects as the basic representation of language semantics. Representations acquired in these conceptual spaces, using the tools of computational geometry, allow for new approaches to generative synthesis of sentences for robot dialog interactions.

Förster et al., in “[Robots Learning to Say ‘No’: Prohibition and Rejective Mechanisms in Acquisition of Linguistic Negation](#),” extend research on symbol grounding to encompass psychological aspects of affect and motivation as they pertain to the process by which a robot acquires negation words. Their experimental study found that negation words are prosodically salient within prohibitive utterances and negative intent interpretations and, thus, can be isolated from a speech signal. The study not only presents an opportunity to develop robot learning approaches, but also to test out some of our theories of child development in the process.

Doering et al., in “[Neural Network-based Memory for a Social Robot: Learning a Memory Model of Human Behavior from Data](#),” build upon the state-of-the-art in machine learning from human-human interaction data by developing a robot memory model that incorporates temporally distant events in decision-making about the best action to take in HRI. Their investigation explores how an automated behavior learning system can learn a memory representation of interaction history with human customers within a simulated camera shop scenario. The results show that a neural network-based model that allows the robot to incorporate memories is more successful at performing appropriate actions than one that does not incorporate memories.

This issue of ACM THRI would not be possible without the substantial efforts given by the authors, reviewers, and editorial team for Representation Learning. The Editors-in-Chief must particularly note the work and contributions of our guest for this issue: Takato Horii of the University of Electro-Communications, Yukie Nagai of the National Institute of Information and Communications Technology, Takayuki Nagai of the University of Electro-Communications, and our first author, Amir Aly of Ritsumeikan University.

Amir Aly

Guest Editor

Odest Chadwicke Jenkins
Selma Sabanovic

Co-Editors-in-Chief