Long-term learning and adaptation in intelligent personal robotics

Adriana Tapus
Interaction Lab, University of Southern California
941 West 37th Place, RTH 407, MC 0781, Los Angeles, CA, USA
tapus@usc.edu

Robotic systems are now capable of social interaction with human users, presenting a new opportunity for providing customizable personal care. Mounting evidence shows that human users respond more readily to robots than to disembodied alternatives such as computer screens, personal digital assistants, and smart phones.

Learning to adapt our daily behavior as a function of different internal and external factors it's a fundamental characteristic of humans. Creating robots capable of exhibiting similar sophisticated capabilities has proven to be a very difficult task. Therefore, providing an engaging and motivating individualized interaction protocol that is adaptable to user personality, preferences, physical and cognitive impairment, and task performance is a challenge in robotics, especially when working with vulnerable user populations, where a careful consideration of the users needs and disabilities is required.

To the best of my knowledge, no work has yet tackled the issue of robot personality and behavior adaptation as a function of user personality and task performance in the assistive human-robot interaction context. In the work described here, I address these issues and propose a reinforcement learning-based approach to robot behavior adaptation. In the learning approach, the robot incrementally adapts its behavior and thus its expressed personality and empathy, attempting to maximize the task performance. The robot's behavior (and therefore personality and empathy) is expressed through multi-modal cues which include: interpersonal distances/proxemics, verbal (e.g., different personality markers in speech such as verbosity, content polarity, and downtoner hedges), para-verbal (e.g., pitch and volume), and non-verbal communication (e.g., posture and gestures), and activity that will allow the robot to be responsive both in terms of temporal and social appropriateness.

I formulated the problem as policy gradient reinforcement learning (PGRL) and developed a learning algorithm that consists of the following steps: (a) parametrization of the behavior; (b) approximation of the gradient of the reward function in the parameter space; and (c) movement towards a local optimum. The main goal of my robot behavior adaptation system is to enable us to optimize on the fly the three main interaction parameters (interaction distance/proxemics, speed, and verbal and paraverbal cues) that define the behavior (and thus personality and empathy) of the therapist robot, so as to adapt it to the users profile and thus improve the users task performance. The learning system changes the robot's personality, expressed through the robot's behavior, in an attempt to maximize the task performance metric.

The first pilot experimental results (with stroke and dementia patients) provided first evidence for the effectiveness of robot behavior adaptation to user personality, preferences, disability, and performance. The results show a novel stroke rehabilitation tool and dementia cognitive attention utensil that provide individualized and appropriately challenging/nurturing therapy style that measurably improves user task performance and attention.

This work demonstrates the promises of intelligent personal robotics, a new research area with large horizons of fascinating and much needed research. This ongoing research is aimed at developing effective embodied assistive systems, and extending my understanding of human social behavior.