Secret Intelligence Service

Room 15

Notes for Discussion:

WHAT MIGHT THE FUTURE HOLD FOR PEOPLE IN THE GLOBAL, CITIZEN-LESS, ROBOTIZED ECONOMY?

(C-I)

R1502062016

These are some of the notes I made for the purpose of background reading which is always useful, but do bear in mind the remit we have which is; to discuss the implications of those future displaced by humanoids, to assess where these people will likely reside geographically, and the likely implications for security. Note also that certain of this material is dated and that a year constitutes a very long time in this.

Introduction

An understanding of how human beings and robots (humanoids) can successfully interact to accomplish specific tasks is crucial in creating more sophisticated robots that may eventually become an integral part of human societies. A social robot needs to be able to learn the preferences and capabilities of the people with whom it interacts so that it can adapt its behaviors for more efficient and friendly interaction. Advances in human-computer interaction technologies have been widely used in improving human-robot interaction (HRI). It is now possible to interact with robots via natural communication means such as speech. In this paper, an innovative approach for HRI via voice-controllable intelligent user interfaces is described. The design and implementation of such interfaces are described. The traditional approaches for human-robot user interface design are explained and the advantages of the proposed approach are presented. The designed intelligent user interface, which learns user preferences and capabilities in time, can be controlled with voice. The system was successfully implemented and tested on a Pioneer 3-AT mobile robot. 20 participants, who were assessed on spatial reasoning ability, directed the robot in spatial navigation tasks to evaluate the effectiveness of the voice control in HRI. Time to complete the task, number of steps, and errors were collected. Results indicated that spatial reasoning ability and voice-control were reliable predictors of efficiency of robot teleoperation. 75% of the subjects with high spatial reasoning ability preferred using voice-control over manual control. The effect of spatial reasoning ability in teleoperation with voice-control was lower compared to that of manual control.

Human beings and robots work more effectively together following cross-training

Swapping of roles improves efficiency as well as robots' confidence and humans' trust.

Spending a day in someone else's shoes can help us to learn what makes them tick. Now the same approach is being used to develop a better understanding between humans and robots, to enable them to work together as a team.

Robots are increasingly being used in the manufacturing industry to perform tasks that bring them into closer contact with humans. But while a great deal of work is being done to ensure robots and humans can operate safely side-by-side, more effort is needed to make robots smart enough to work effectively with people, says Julie Shah, an assistant professor of aeronautics and astronautics at MIT and head of the Interactive Robotics Group in the Computer Science and Artificial Intelligence Laboratory (CSAIL).

"People aren't robots, they don't do things the same way every single time," Shah says. "And so there is a mismatch between the way we program robots to perform tasks in exactly the same way each time and what we need them to do if they are going to work in concert with people."

Most existing research into making robots better team players is based on the concept of interactive reward, in which a human trainer gives a positive or negative response each time a robot performs a task.

However, human studies carried out by the military have shown that simply telling people they have done well or badly at a task is a very inefficient method of encouraging them to work well as a team.

So Shah and PhD student Stefanos Nikolaidis began to investigate whether techniques that have been shown to work well in training people could also be applied to mixed teams of humans and robots. One such technique, known as cross-training, sees team members swap roles with each other on given days. “This allows people to form a better idea of how their role affects their partner and how their partner’s role affects them,” Shah says.

In a paper to be presented at the International Conference on Human-Robot Interaction in Tokyo in March, Shah and Nikolaidis will present the results of experiments they carried out with a mixed group of humans and robots, demonstrating that cross-training is an extremely effective team-building tool.

To allow robots to take part in the cross-training experiments, the pair first had to design a new algorithm to allow the devices to learn from their role-swapping experiences. So they modified existing reinforcement-learning algorithms to allow the robots to take in not only information from positive and negative rewards, but also information gained through demonstration. In this way, by watching their human counterparts switch roles to carry out their work, the robots were able to learn how the humans wanted them to perform the same task.

Each human-robot team then carried out a simulated task in a virtual environment, with half of the teams using the conventional interactive reward approach, and half using the cross-training technique of switching roles halfway through the session. Once the teams had completed this virtual training session, they were asked to carry out the task in the real world, but this time sticking to their own designated roles.

Shah and Nikolaidis found that the period in which human and robot were working at the same time — known as concurrent motion — increased by 71 percent in teams that had taken part in cross-training, compared to the interactive reward teams. They also found that the amount of time the humans spent doing nothing — while waiting for the robot to complete a stage of the task, for example — decreased by 41 percent.

What’s more, when the pair studied the robots themselves, they found that the learning algorithms recorded a much lower level of uncertainty about what their human teammate was likely to do next — a measure known as the entropy level — if they had been through cross-training.

Finally, when responding to a questionnaire after the experiment, human participants in cross-training were far more likely to say the robot had carried out the task according to their preferences than those in the reward-only group, and reported greater levels of trust in their robotic teammate. “This is the first evidence that human-robot teamwork is improved when a human and robot train together by switching roles, in a manner similar to effective human team training practices,” Nikolaidis says.

Shah believes this improvement in team performance could be due to the greater involvement of both parties in the cross-training process. “When the person trains the robot through reward it is one-way: The person says ‘good robot’ or the person says ‘bad robot,’ and it’s a very one-way passage of information,” Shah says. “But when you switch roles the person is better able to adapt to the robot’s capabilities and learn what it is likely to do, and so we think that it is adaptation on the person’s side that results in a better team performance.”

The work shows that strategies that are successful in improving interaction among humans can often do the same for humans and robots, says Kerstin Dautenhahn, a professor of artificial intelligence at the University of Hertfordshire in the U.K. “People easily attribute human characteristics to a robot and treat it socially, so it is not entirely surprising that this transfer from the human-human domain to the human-robot domain not only made the teamwork more efficient, but also enhanced the experience for the participants, in terms of trusting the robot,” Dautenhahn says.

Cognitive Robotics and Human Robot Interaction

This approach to Human-Robot Interaction is through cognitive robotics: understanding how and why people act the way they do. More capable and intelligent robots and autonomous systems will require more human-like cognitive abilities.

The hypothesis is that robots and autonomous systems that use human-like representations, strategies, and knowledge will enable better collaboration and interaction with the people who use them. Similar representations and reasoning mechanisms make it easier for people to work with these autonomous systems. An autonomous system must be able to explain its decisions in a way that people understand, which should lead to better trust and acceptance of the system. If an autonomous system can predict a person's needs, even in the very short term, it can prepare for it and act appropriately.

In this line of research, *computational cognitive models* are used to build process models of human cognitive skills, and those models are then used as as reasoning mechanisms on the robots and autonomous systems. We build computational cognitive models of people -- their perception, their memory, their attention, their reasoning, their spatial abilities, and their thinking. We use an embodied version of ACT-R (Anderson et al., 2007) that we call ACT-R/E (Trafton et al., 2013). ACT-R (and ACT-R/E) are computational systems that are based on theories of how human reasoning work, and which capture known facts and constraints known about how the mind works, and connect well with psychological data (experiments) and neuroscience data (fMRI).

There are two primary scientific goals:

- To understand the embodied nature of cognition: how people work in the physical world.
- To improve human robot interaction by high fidelity models of individuals so that we can provide some assistance to them. For example, our models understand that people do not have perfect memories and cannot see behind their head. This knowledge allows our model to remind a person what they were doing if they forgot or to show them something in the environment they didn't see.

Some of the cognitive models that have been developed and have been used in various research projects include:

- Gaze following: The ability of an infant around the age of 18 months to follow objects in the environment (such as a toy).
- Level 1 Perspective Taking: The ability to understand what another person is pointing at that developed around the age of two years.
- Visual, Spatial Perspective taking via mental simulation: Around the age of 4-5 years of age, a child can mentally simulate how the world looks from someone else's point of view.
- Conversation tracking: Being able to follow several people engaged in conversation and knowing where to look when during conversations.

- Teaming via model of one's self: Allows deciding what a team mate will do based on modeling the team mate as one's self.
- Theory of Mind: The ability to infer the beliefs, desires and intentions of others, which develops around the age of 5

Do Not Read my Face
Tackling the Challenges of Facial Masking in Parkinson's Disease Rehabilitation through Co-Robot Mediator

Many aspects of co-robots are currently investigated, from algorithms for scene and activity understanding, to planning for human-robot teaming, and natural language interactions between humans and robots. However, there is surprisingly little work on mechanisms that will allow co-robots to behave in a manner that is ethical and sensitive to the moral context and social norms. This is particularly worrisome as simple robots are already entering society without any notion of ethically acceptable behavior and this situation will only be exacerbated in the future if various kinds of social and assistive robots will cause humans to form unidirectional emotional bonds with robots without those robots being sensitive to human emotions and feelings. In this NSF-funded collaborative NRI project, we will tackle a hitherto completely overlooked ethical aspect of human-robot interaction: *maintenance of human dignity and the stigmatization of human patients*.

The overarching scientific goal of this project is two-fold: (1) to develop a robotic architecture endowed with moral emotional control mechanisms, abstract moral reasoning, and a theory of mind that allow co-robots to be sensitive to human affective and ethical demands, and (2) to develop a specific instance of the architecture for a co-robot mediator between people with "facial masking" due to Parkinson's disease (PD) that reduces their ability to signal emotion, pain, personality and intentions to their family caregivers, and health care providers who often misinterpret the lack of emotional expressions as disinterest and an inability to adhere to treatment regimen, resulting in stigmatization. Specific questions we will address include:

(1) How can an expanded set of moral emotions, particularly empathy, be modelled and exhibited by co-robots to provide quantitatively better care of patients, in particular, early

patients with PD.

(2) How can we develop a theory of mind of both caregiver and patient (including their goals and emotional states) that can be used by a “co-robot mediator” to improve the quality of care for patients while enhancing the dignity of both patient and caregiver?

To tackle these problems, the project brings together two roboticists, Prof. Matthias Scheutz (Tufts) and Prof. Ron Arkin (Georgia Tech) with extensive prior experience in robot ethics and modeling emotions as well as implementing them in integrated autonomous robotic systems. The robotics expertise is combined with that of an expert in early PD rehabilitation and daily social life, Prof. Linda Tickle-Degnen (Tufts).

(C-I)

Room 15