Learning, Planning and Sharing Robot Knowledge for Human-Robot Interaction

Summary of Dagstuhl Seminar 10401, October 3–8 2010

Rachid Alami (LAAS - Toulouse, FR)
Rüdiger Dillmann (KIT - Karlsruhe Institute of Technology, DE)
Thomas C. Henderson (University of Utah, US)
Alexandra Kirsch (TU München, DE)

January 10, 2011

Autonomous robots acting in unstructured environments need sophisticated cognitive capabilities including perception, manipulation and navigation, planning and reasoning as well as knowledge acquisition and processing skills. Only when all these capabilities are combined can the robot show intelligent behavior and act appropriately in environments that are primarily designed for humans. Given the growing availability and inter-connectivity of modern agents and robots, better mechanisms to define, learn and share knowledge must be developed. Advances in robot platforms have led to their integration into society for a variety of functions, and there is a pressing need to understand how they can acquire and exploit knowledge required for their specific activities. In order to be useful, robotic agents must be able to recognize a wide range of objects, relations, and situations in their environment, and to understand the semantics of these. In addition, it is imperative that procedural and process knowledge also be known so that interactions are safe and meaningful. Cooperation and collaboration are also essential as well as the ability to perceive human and animal emotions and intentions to the largest degree possible.

This seminar brought together 25 scientists with experience in planning, learning, knowledge processing and human-robot interaction to discuss how the high-level control concepts are to be used and adapted to human-robot interaction. Because high-level control of robots interacting with humans is of special practical importance, we invite experts with experience in assistive technology and those examining social acceptance of such systems.

The seminar consisted of 8 sessions, which were focussed on discussions of the topics knowledge sharing, learning and interaction. In addition, there were 3 short presentations on topics that arose in the discussions and 7 full presentations. In the following we summarize the group discussions.

Knowledge sharing. How can robots share knowledge among themselves and how can knowledge be shared between robots and humans?

For sharing knowledge among robots the WWW seems a very promising approach, together with specific search engines that facilitate the recovery of specific knowledge and skills that a robot needs to fulfill a given task.

A very basic level of knowledge sharing is the exchange of sensor data acquired by different robots. Sharing the raw data has the advantage that no intermediate representation levels

Dagstuhl Seminar Proceedings 10401
Learning, Planning and Sharing Robot Knowledge for Human-Robot Interaction
http://drops.dagstuhl.de/opus/volltexte/2011/2934

(and with this a common semantic) are necessary and that all information contained in the original data is passed on. On the other hand, even such a simple form of representation would require meta-information on how the data was acquired and under which circumstances to make it useful for other robots.

A specific way of sharing procedural knowledge is the availability of robot skills on the WWW. This would allow robots to download skills that they need to fulfill a certain task, relying on the experience of other robots. An interesting question in this context is if the download of a skill should be considered as a form of learning.

Sharing knowledge between robots and humans requires more sophisticated levels of representation and must be embedded into some form of dialogue between humans and robots (see also paragraph on non-physical interaction).

Learning. One fundamental question that was discussed during the seminar is: What is learning? In the discussion it became clear that there are lots of different aspects that a robot needs to learn and that each problem requires specific learning approaches. The following categories of learning were identified: Exploration, Observation (diadic, triadic), Interaction & Teaching, Adaptation & Optimization.

The question why a robot should learn at all was also discussed. In the discussion several points were mentioned:

- "Efficiency", "Intelligence for cost-down"
  By learning skills can be transfered to different (hardware) platforms and the design of robot behavior could be facilitated by a general learning architecture and the robot's own ability to learn the necessary skills and knowledge.
- Solve problems that humans can't currently solve explicitly by learning
  For example language acquisition and transfer of implicit skills to machines
- Acquire information not existing at design time
- Acceptance in Human Robot Interaction
  Learning is perceived as very human, hence learning machines will be better accepted

The discussion also included architectures and the relation between higher-level abstract reasoning processes and lower-level navigation and manipulation actions (Figure 1). Learning is necessary on all levels of abstraction, but the interaction between these levels is not yet clear.

Physical Interaction. There are specific challenges when robots interact physically with a human, also including aspects of safety.

The sensors available for robots are still far from the capabilities of humans. Although touch sensors and artificial skins are becoming available, the resolution is far from that of the human skin. Besides, more sensor information raises the problem of processing these large amounts of data.

As an example for the discussion, the question was raised what is missing in the state of the art to implement a robot massaging people (with comparable skill as human massagers). There are already automatic massage devices, which are however far from the abilities of humans. A massage robot would require very accurate sensing capabilities and would have to

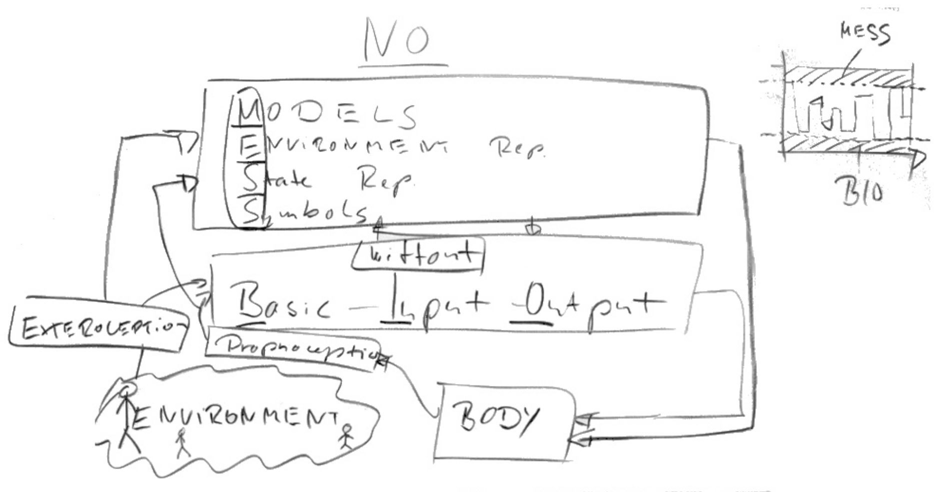

Figure 1: Agent architecture as a basis for learning.

adapt its behavior for each individual unlike the machines of today, which perform the same sequence of movements for all users.

Non-physical Interaction. Robots also interact with people in non-physical ways, for example for communication.

Speech is often used as a natural communication modality. While speech recognition works reasonably well for well-defined vocabularies, a general solution is still out of reach. Beside the capabilities of understanding spoken language, it is very important not to create unrealistic expectations when designing robots. For example, a robot that uses speech output is expected also to understand spoken language. Especially when the robot uses the word "I" humans interpret something like a personality for the robot and hence expect it to behave like a person. On the other hand, if a robot doesn't give the impression to have a certain capability, people won't expect it (e.g. a robot without arms is not expected to manipulate objects).

There is also the question of which communication modalities people want to use when interacting with robots and how often they wish to communicate. Studies have found that users like speech interaction in private domains, but in public (e.g. for vending machines) they prefer not to speak to machines.

Another question for interaction is how to keep humans engaged when interacting with a robot. There was a discussion if this question should not rather be solved by providing useful functionality to the robot rather than designing specific interation behavior that keeps humans interested in the robot. It seems that both aspects are important for the acceptance of robots.