AISB Quarterly, 2003, Vol. 113, pp. 6-7.

INSTRUCTING ROBOTS

Guido Bugmann
Robotic Intelligence Laboratory
University of Plymouth

Knowledge transfer between a user and a new robot assistant is necessary to make the robot functional. In the same way, human employees need instruction before being able to fulfil their duties. In industrial robotics, robots are programmed by a small number of trained operators. In domestic or service robotics, users are in large number and cannot be expected to know or learn a programming language. Instead, future robot assistant may need to be able to understand instructions issued in natural language and to convert them into corresponding programs' code. Therefore, a joint EPSRC project between the Universities of Plymouth and Edinburgh was aimed at designing a system for programming robots by using spoken input and identifying the limitations of current tools.

In the "Instruction Based Learning (IBL)" project, a mobile robot is instructed how to travel from one place to another in a miniature town (see figure 1 and ref. IBL). On the basis of its instructions, it then creates a computer program that it uses to navigate between the two places. The complete system, starting with spoken input and ending with a navigating robot, has been built and was recently exhibited in the Plymouth City Museum as part of a robotics week. In that noisy environment, speech recognition proved surprisingly robust, failing mainly on the high-pitched voice of some children. Indeed, there are still unsolved problems as outlined in following sections.

Figure 1. Experimental set-up. A miniature remote-brained robot with a 8cm x 8cm base performs vision-guided navigation in a 170 cm x 120 cm model town, following a sequence of actions defined in prior instructions given by a user. Video images are sent by wireless from the robot to a PC for processing. The resulting motion commands are sent back by wireless to the robot. The user speaks to the robot through a headset microphone connected to the PC.

The first step in the design of the IBL system was the recording of a corpus of 144 route instructions given by 24 subjects. They were told to speak for a human operator who would later move the robot by remote control. The transcripts of the utterances (6600 words, 330 distinct words) were used to select, out of a wide-coverage grammar, the restricted set of grammatical rules and lexicon corresponding to this domain. This user-centred approach to the design of restricted grammars enables in principle users to use unconstrained speech, while maximizing speech recognition performance. The results however showed that wide-coverage grammars do not cover some of the forms found in spoken language, and 40% of utterances were not covered. Further, there were indications that the domain lexicon was not closed and that about one new word was to be expected for every two new instructions (Bugmann et al., 2001). Both problems do not have accepted solutions yet and further research is required, for instance towards the development of grammars for spoken language and mechanisms for dealing with out-of-grammar words.

The corpus underwent a further functional analysis to determine which navigation actions users refer to in route instructions. A list of 14 primitive functions was established including functions such as “turn in direction x after the nth landmark y” or “the goal x is located in relation y to landmark z” (Lauria et al., 2002a). Each of these functions was pre-programmed into the system and a new procedure specification language (PSL) was created to encode rules that map natural language expressions to the appropriate function calls (Lauria et al., 2002b). There are multiple ways to refer to a given action and about 200 rules were required. Again, the list of primitive functions was found not to be closed, with about one new function expected for every 35 new instructions. It is unclear at present how this problem should be dealt with, because a primitive is a piece of low-level robot program that the user cannot create. Here, methods of learning by example may prove useful (see e.g. Schaal et al., 2003).

The combination of primitives defined in the instruction is converted into a new piece of program code (using the scripting language PYTHON) having the same access protocol as pre-programmed primitives. Thus, learnt procedures can be re-used in later instructions to create more complex procedures (figure 2). However, natural language references to previously taught procedures revealed a range of new problems only partially solved in the IBL project (Lauria et al, 2002b).

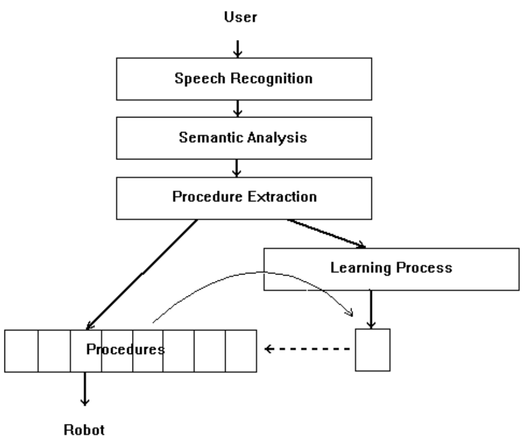

Figure 2. IBL concept diagram. A command issued by the user can either be executed immediately if the robot has the corresponding procedure, or a learning process is initiated. Thereby, a new procedure is created from a combination of existing primitives.

Primitives are actions referred to in human-to-human instructions. They correspond to common human execution capabilities such as finding a left turn and taking it. For robots, however, these are not simple functions and it is not straightforward to write a turn() function that can deal with every type of left turn (Kyriacou et al., 2002). Thus, a robot that understands unconstrained natural language instructions needs quite advanced action capabilities. Or, conversely, a robot with limited action capabilities can only understand a limited range of expressions. As future users cannot be expected to know all the robot's capabilities, an information mechanism needs to be built in the robot. Interestingly, this turns out to be the same problem as teaching the user which expressions the robot can understand. Future research along these lines may also solve many of the natural language processing problems above, as the user would seek to adapt his/her language to that of the robot. Demonstrations have suggested that this is a natural tendency in users.

Acknowledgments:

Following researchers contributed to the IBL project: Ewan Klein, Johan Bos, Stanislao Lauria, Theocharis Kyriacou and Kenny Coventry.

References:

G. Bugmann, S. Lauria, T. Kyriacou, E. Klein, J. Bos and K. Coventry (2001), Using Verbal Instruction for Route Learning: Instruction Analysis, Proc. TIMR 01 – Towards Intelligent Mobile Robots, Manchester 2001. Technical Report Series, Department of Computer Science, Manchester University, ISSN 1361 – 6161. Report number UMC-01-4-1.

IBL: http://www.tech.plym.ac.uk/soc/staff/guidbugm/ibl

T. Kyriacou, G. Bugmann, S. Lauria (2002), Vision-Based Urban Navigation Procedures for Verbally Instructed Robots, Proceedings of the 2002 IEEE/RSJ Intl. Conference on Intelligent Robots and Systems (IROS'02) EPFL, Lausanne, Switzerland • October 2002. pp.1326-1331.

S. Lauria, G. Bugmann, T. Kyriacou and E. Klein (2002a), Mobile Robot Programming Using Natural Language, Robotics and Autonomous Systems, 38 (3-4): 171-181.

S. Lauria T. Kyriacou G. Bugmann, J. Bos, and E. Klein (2002b), Converting Natural Language Route Instructions into Robot-Executable Procedures, Proceedings of the 2002 IEEE Int. Workshop on Robot and Human Interactive Communication (Roman'02), Berlin, Germany, pp. 223-228.

S. Schaal, A.J. Ijspeert and A. Billard, A. (2003), Computational Approaches to Motor Learning by Imitation, Philosophical Transactions: Biological Sciences (The Royal Society), 358:1431, p.537-547.