Intelligent Robotic Systems
Editorial Introduction

Ronald C. Arkin
Mobile Robot Laboratory
College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280

1. Introduction

Just what are Intelligent Robotic Systems, anyway? Our working definition for an intelligent robot has been:

> A machine that is able to extract information from its environment and use knowledge about its world to move safely in a meaningful and purposeful manner.

There are others that would add to this definition. Schank requires that learning be a requisite part of any artificially intelligent system [18]. Arbib [2] adds world model maintenance and planning capabilities. While it is agreed that these are potentially useful additions, they are not defining characteristics. It is rather the active participation of a robotic agent with its world (not just a simulation) that makes a robot a robot.

This moves a large burden to the perceptual and motor control aspects of the robotic system, an area often neglected by those in mainstream artificial intelligence. Indeed it is our premise that without a thorough understanding of a robotic agent's ecological relationships, it is impossible to effectively construct an intelligent robotic system. The notion of thus viewing a robot as inseparable from its environment is central to this position [5] and is consistent with views from a biological/psychological perspective (e.g., [11,13,12]).

2. No Excuses

In order to effectively study robotic agents then, it becomes necessary to ground them in real situations, using physical agents whenever and wherever possible. This requirement is reflected in **all** of the papers included in this special issue. Intelligent robotics requires real world robots. Simulations, although useful as stepping stones for testing and exercising software, are not enough.

Fortunately, there are no longer any valid excuses for researchers not to use real robots. These systems, which were exorbitantly expensive as near as a decade ago, are now affordable and even within the reach of hobbyists with a bent for hardware. The AAAI conference now annually hosts a competition where robots are actually constructed and programmed on-site by inexperienced users. Thus with the burgeoning number of options for physical robots, beginning with systems at very low-cost, reviewers and editors are now able to insist that theories be fielded on working systems (or at least have a clear path to such implementation) prior to their publication in archival journals. This real-world filter will help separate potentially fallacious theories from those that are truly valid. Robotics is an experimental science and must have realistic experiments as a central component for verifying hypotheses. The general scientific method can and should be applied to our field.

Available hardware is only expected to improve. There still, however, are many sensors that are unavailable to many researchers, either in terms of cost or size (laser scanners are one such component). Fortunately, costs are continuing to fall on much of this technology while performance improves. This parallels, to a lesser degree, the well-known decrease in computational hardware costs, which now also makes most basic computer vision research affordable. Newer technologies, such as micromachines, also promise to revolutionize the field. Progress is, unfortunately, not being made equally on all of the remaining problems: bottlenecks still exist with lightweight durable materials, low cost high-endurance power supplies, and actuation at a distance, among others.

3. Architectural Issues

An important conceptual development involves the revolution of organizing principles by which intelligent robotic systems are structured: robot architectures. Much has transpired from the early resolution-theorem proving days of Shakey [16]. Traditional architectures, those with a heavy reliance on symbol processing and map construction were the early mainstay of robotic systems which eventually led to the generation of an official governmental standard, the hierarchical NASREM architecture [1] in the 1980s.

Actual progress, however, seemed to be move at a snail's pace (whether hardware limited or architecturally limited is a debatable point). This led to the abandonment of many of the feaures of the hierarchical model in favor of a reactive, behavior-based strategy that emerged in the mid-1980s [6,3,14]. This, was a radical change in direction; from top-down human-oriented models, to bottom-up low-level biological models (e.g., ants, frogs, and cockroaches). The net result was the appearance of robots more capable of participating in dynamic environments than their predecessors. These systems did not aspire to human-level performance: that was not the point [7]. The intent was rather to create task-oriented robotic systems that played specific roles within their world and carried out those roles well. Other recent efforts have been concerned with hybridizing reactive and hierarchical systems (e.g., [4,9]) and introducing effective learning capabilities.

4. This issue

All four papers in this special issue reflect to some degree these architectural considerations. Grupen et al's paper is an example of a reactive system which provides distributed reusable control modules. This is illustrated in a particularly challenging domain, control of a hand/arm robotic system with 20 degrees of freedom.

Montgomery, Fagg, and Bekey's contribution describes the first implementation of a behavior-based architecture on a real flying robot. Loosely coupled control modules are used within the confines of a hierarchical organization. Testing these ideas on a helicopter is a particularly challenging task due to the unforgiving nature of crashes.

Liscano et al's paper is illustrative of a hybrid system incorporating aspects of both reactive and hierarchical control. Using a blackboard architecture to support multiple activities, the system was testing on a modified Cybermotion platform for typical navigational tasks using range data.

Finally, Kaiser et al's paper presents an introduction of learning capabilities to these systems. Learning and adaptation in real robotic systems represents an important marriage between robotics and the machine learning community. A recent special issue on the subject confirms this [17]. Kaiser's research, reported in this issue, continues that tradition, with the additional benefit of experimental results from two working robotic vehicles.

5. Moving to the Big Time

The results of a decade of progress are being realized in large scale intelligent robotic systems now being fielded. Several representative examples are listed below:

- SWAMI II Stored Waste Autonomous Mobile Inspector [15], is a robot scheduled to conduct low-level nuclear waste surveillance operations in a Department of Energy site, beginning in the summer of 1995.

- DANTE II [20] this past summer successfully met its scientific goals by descending into an Alaskan volcano and carrying out experiments in a location forbidden to human entry. Although the system had difficulty on its return, it nonetheless constitutes a major advance in fielded working systems.

- A microrover [10] slated to be launched in 1996 on NASA's Pathfinder mission to Mars to conduct simple experiments within the area surrounding the lander.

- ARPA, as part of a series of technology demonstrations called the UGV Demo II program [8], is developing a team of four Jeep-like robots capable of conducting military scout missions in rugged outdoor terrain, with the final demonstration slated for the summer of 1996.

- The Army Applied Aviation Directorate is sponsoring the development of Autonomous Scout Rotorcraft Testbed (ASRT) which will serve as an autonomous helicopter scout, to be demonstrated in the summer of 1996 in Atlanta.

6. What next?

From these examples, one can easily see that robotic technology is on the verge of having real impact on real world problems. Research on intelligent robotic systems, of the type reported in this special issue, has led us to this point. Optimistically, the age of robots is beginning to arrive. Nonetheless, there will be significant growing pains in these early years and patience will continue to be required if the efforts of robotics researchers are to bear fruit. Many outstanding questions remain such as:

- Just what is the appropriate role of representational knowledge within intelligent robotic systems?
- What are the most effective methods for the integration of hierarchical, reactive, and/or neural architectures?
- How can continuous adaptation and evolution for ecological niche finding in robotic systems be performed? (Sims has performed some very interesting pioneering work in this area [19])
- How can spatially distributed perception/actuation be coordinated among teams of robots?

This is merely a small subset of the largely unanswered questions confronting roboticists. There certainly is no shortage of work to be done.

Acknowledgments

I would like to acknowledge those who made this special issue possible. B. Chandrasekharan and Steve Cross, past and present editors-in-chief of IEEE Expert respectively, encouraged and enabled this effort during all phases from inception to completion: I an indebted to them for their assistance. The contributors' cooperation, responsiveness to the reviewer's comments, and their timely delivery of manuscripts has been greatly appreciated. Finally, I'd like to express deep appreciation to the reviewers themselves for the time and effort they donated in ensuring the quality of this issue.

REFERENCES

[1] Albus, J., McCain, H. and Lumia, R., “NASA/NBS Standard Reference Model for Telerobot Control System Architecture (NASREM)”, NBS Technical Note 1235, Robot Systems Division, National Bureau of Standards. 1987.

[2] Arbib, M., *The Metaphorical Brain*, Wiley-Interscience, New York, 1972.

[3] Arkin, R.C., “Motor Schema Based Navigation for a Mobile Robot: An Approach to Programming by Behavior”, *Proceedings of the 1987 IEEE Conference on Robotics and Automation*, Raleigh, NC, pp. 264-271.

[4] Arkin, R.C., “Integrating Behavioral, Perceptual, and World Knowledge in Reactive Navigation”, *Robotics and Autonomous Systems*, 6 (1990), pp. 105-122.

[5] “Just What is a Robot Architecture Anyway? Organizing Principles versus Turing Equivalency”, to appear in working notes of *AAAI Spring Symposium on Lessons Learned from Implemented Software Architectures for Physical Agents*, March 1995, Stanford, CA.

[6] Brooks, R, “A Robust Layered Control System for a Mobile Robot”. *IEEE Jour. of Robotics and Automation*, 2(1):14-23, 1986.

[7] Brooks, R., “Elephants Don’t Play Chess”, *Robotics and Autonomous Systems*, 6 (1990), pp. 3-15.

[8] Chun, W. and Jochem, T., “Unmanned Ground Vehicle Demo II: Demonstration A”, *Unmanned Systems*, Vol. 12, No. 1, Winter 1994, pp. 14-20.

[9] Gat, E., “Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots”, /it Proc. AAAI-92. Morgan-Kaufmann.

[10] Gat, E., Behar, A, Desai, R., Ivlev, R., Loch, J. and Miller, D.P., “Behavior Control for Planetary Exploration: Interim Report”, *Proc. IEEE International Conference on Robotics and Automation*, May 1993, pp. 567-71.

[11] Gibson, J.J., “The Theory of Affordances”, *Perceiving, Acting, and Knowing*, ed. Shaw, R., and Bransford, J., Erlbaum, 1977.

[12] Neisser, U., *Cognition and Reality: Principles and Implications of Cognitive Psychology*, Freeman, 1976.

[13] Norman, D. and Shallice, T., “Attention to Action: Willed and Automatic Control of Behavior”, from *Consciousness and self-regulation: Advances in research and theory*, Vol. 4, R. Davidson, G. Schwartz, and D. Shapiro, (eds)., Plenum Press, pp. 1-17, 1986.

[14] Payton, D., “An Architecture for Reflexive Autonomous Vehicle Control”, *Proc. IEEE Conf. on Robotics and Auto.*, New York: IEEE, April 1986, pp. 1838-1845.

[15] Peterson, K. and Ward, C., “An Autonomous Mobile robot to Perform Waste Drum Inspections”, *Robotics and Manufacturing - Recent Trends in Research, Education, and Applications*, Volume 5, ed. M. Jamshidi et al, ASME Press, 1994, pp. 637-641.

[16] Nilsson, N., ed., (1984). “Shakey the Robot”, SRI International *Technical Note 323*.

[17] *Robotics and Autonomous Systems*, Special Issue on Robot Learning, Vol. 8, No. 1-2.

[18] Schank, R.C., “What is AI anyway?” *AI Magazine*, Winter 1987, pp. 59-65.

[19] Sims, K., "Evolving Virtual Creatures", *Proc. SIGGRAPH 94*, July 1994, pp. 15-22.

[20] Wettergreen D., Thorpe, C., and Whittaker, R., "Exploring Mount Erebus by Walking Robot", *Robotics and Autonomous Systems*, Vol. 11, no. 3-4 Dec. 1993, pp. 171-185. (This article describes Dante, the predecessor to the more successful Dante II mission).