



# Semisentient Robots: Routes to Integrated Intelligence

Luís Seabra Lopes, *Universidade de Aveiro*  
Jonathan H. Connell, *IBM T.J. Watson Research Center*

Complex systems are getting to the point where it almost feels as if “someone” is there behind the interface. This impression comes across most strongly in the field of robotics because these agents are physically embodied, much as humans are. We believe that this phenomenon has four primary components: A system must be able to

- act in some reasonably complicated domain,
- communicate with humans using a language-like modality,
- reason about its actions at some level so that it has something to discuss, and
- learn and adapt to some extent on the basis of human feedback.

These components all contribute to various aspects of the definition of “sentience.” Yet, we obviously can combine these basic ingredients in different ways and proportions. This special issue, therefore, examines what sort of interesting recipes people are cooking up.

## Perception & action

Philosophers have studied the nature of intelligence since ancient times. In modern times, researchers from other fields, including psychology, cognitive science, and AI, have joined the debate. Unfortunately, while we can define some abilities related to intelligence, such as learning or reasoning, with some precision, the word *intelligence* means different things to different people. Humans are intelligent, everybody seems to agree. But are animals intelligent? Is a computer running a medical expert system intelligent? In the larger picture, is sentience even possible for a silicon-based entity such as HAL 9000?<sup>1</sup>

One precondition for intelligence seems to be a responsiveness to external stimuli. That is, an agent must perceive something about its environment and be able to perform appropriate actions as conditions change. Let’s call this the *animate* criterion. This does not mean an agent must have a camera or a robotic arm, merely that it has some means of input and output. This might be as simple as a text-based system. The point is that an agent should not always emit the exact same response; at the least, some changes in its input should lead to changes in its output. Certainly we would not consider intelligent a robot that just sat on a table or drove forward. Given that perception and action are definitely required, determining such a system’s intelligence comes down to judging its actions’ appropriateness in light of the prevailing conditions. In fact, robotics is often defined as “the intelligent connection of perception to action.”<sup>2</sup> The question now becomes, how do we make this connection?

The 1950s and 1960s saw the creation of the first robotic devices, such as Elsie<sup>3</sup> and the Hopkins Beast.<sup>4</sup> These were inspired by biological models of simple creatures, akin to the current A-life and animat approaches. These machines did things such as react to lighting variations and follow along hallways looking for “food.” Although they definitely had a quality of aliveness, they did not seem particularly intelligent in the human sense.

In the late 1960s, the first robots showing some degree of human-like intelligence and autonomy appeared. The mobile robot Shakey, developed at SRI, is generally considered a watershed in autonomous systems. Shakey could already perform some functions that we still consider fundamental for autonomous systems—namely, perception, representation, planning, execution monitoring, learning, and failure recovery. One big advance was that the system was symbol-based, so a human could clearly understand what the system was doing and why. At about the same time, Terry Winograd created SHRDLU, a natural-language-understanding system for a physically embedded agent, a simulated robot arm manipulating blocks.<sup>5</sup> Again, this system seemed more intelligent because it could understand typed symbolic directions much as a human would.

Unfortunately, these classical AI techniques and representations did not extend easily to complex dynamic environments.<sup>6–8</sup> Researchers had difficulty maintaining world models and connecting the symbols in these models to numerical variables derived from the outside world. Furthermore, AI systems tended to be very slow, because of the sequential processing of information and the cost of maintaining the world model. So, they had difficulty reacting to sudden changes in the environment. This is partially related to the combinatorial-explosion problem in planning, which also limits a system's decision-making capabilities. Finally, these early AI systems were brittle; they failed in situations only slightly different from those for which they were programmed. They assumed their world models were correct and complete and thus tended to run blindly, in a ballistic fashion.

These realizations sparked new interest in the animal-modeling paradigm, but now leavened with insights from traditional computer science. Probably the most influential of these new *behavior-based* approaches was Brooks' *subsumption architecture*.<sup>8</sup> This approach supports competence levels that range from basic reflexes (reactions) all the way to reasoning and planning. The more abstract control layers exploit the functionality implemented by the lower layers but can also directly access sensor data and send actuator commands. Each layer is implemented by a set of modules that run in parallel, without any centralized control.

In such a system, a centralized world model is not necessary (or possible). Instead, the system substitutes continual sensing for

potentially faulty memory, following the maxim that “the world is its own best model.” This decentralization can go even further by completely removing all internal representation and reasoning from the robot control architecture: “intelligence without reason.” The extent to which these robots appear intelligent is largely in the mind of the beholder—such intelligence is not explicitly encoded. Rather, the robot's functionality emerges dynamically from the interaction among its components and the environment.<sup>9</sup>

Reactive and behavior-based systems provide a robust solution to the problem of situated activity, coping with dynamically changing environments. In particular, one of the approach's successes is the fact that these reactively controlled robots could display emergent intelligent-looking behavior. However, despite the modular and incremental design that reactive systems often use, building them seems difficult, requiring considerable engineering effort to hard-code the desired functionality. This is due partly to unexpected interactions between behaviors as the system's size scales up.

In addition, the representationless character of behavior-based systems is often criticized as too restrictive. For instance, mobile-robot path planning requires some representation of space and spatial relations between objects, which is most conveniently conceptualized in a centralized framework. Moreover, with massive decentralization, dynamically imparting goals to such systems is difficult because there is no centralized “it” to give the commands to. Of course, researchers have devised clever ways around some of these restrictions.<sup>10,11</sup>

Nevertheless, while some researchers continue to investigate purely reactive systems, many others have proposed hybrid systems.<sup>7,12–17</sup> Typically, such systems employ reactivity in low-level control and employ explicit reasoning for higher-level deliberation.

## Reasoning

One way to make an embodied system appear more intelligent is to give it the ability to reason. By this we mean the process of stringing together a number of small inference or planning steps to achieve some goal or reach some conclusion. The analogy with humans suggests that robots should be able to plan sequences of actions for achieving given goals and should be prepared to handle unforeseen situations (exceptions) that occur when executing those actions. Plan-

ning (either for a given task or for exception handling), execution monitoring, and explanation (diagnosis) of exceptions are key reasoning functions for robots that will provide a high-level interface to humans.<sup>17,18</sup>

By providing the robot a library of primitive routines and a suitable composition methodology, we can potentially cover a large variety of situations with a relatively modest amount of knowledge. Indeed, this can even give the robot the flexibility to handle additional environmental challenges that we hadn't explicitly considered when developing the library of primitives. The robot does not necessarily need a canned operator to handle every contingency—consider the wide variety of chess board configurations that can result over time from the relatively simple moves of the pieces. In this way, reasoning makes the robot more adaptable.

Most current planning and reasoning systems are derived from either *first-order predicate logic* or the *situation calculus*.<sup>19</sup> The FOPL approach employs logical formulas that describe the world's current state, classify objects, and predict the effects of actions. Often, in robotics, it is assumed that the sensor systems continuously update the truth values of those formulas. The situation calculus and its more modern cousin, Golog,<sup>20</sup> deal primarily with action planning. Each operator (action) has a precondition list governing its applicability to a given situation, and a postcondition list describing its effect on the environment in terms of added and dropped assertions. Usually, a set of frame axioms describe what happens to assertions not directly referenced in the postcondition lists. Both types of systems are attractive because complete and correct proof procedures exist. Also, some hope exists that programming can be purely declarative—the user specifies a goal or points out some nonoptimal condition, and the robot automatically determines how to achieve the desired result.

A variety of complementary model-based techniques, such as qualitative reasoning, diagnostic reasoning, and probabilistic reasoning, are also under investigation. Because these are potentially relevant for building truly sentient robots, they might gradually make their way into future robots.

The *case-based approach* to reasoning is radically different.<sup>21</sup> This approach matches the current situation as the robot perceives it against all the prototypes in its library. Typically, the closest match above some mini-

mum threshold becomes the model for guiding subsequent interactions with the environment. Advanced systems sometimes attempt to combine several known cases to generate better solutions. By design, these systems are good at handling imprecise specifications. However, their case-matching metrics must be carefully crafted. Moreover, these systems are prone to hallucinate and sometimes apply inappropriate models when the available information is sparse. So, case-based systems, although remaining symbol-based and escaping some of the brittleness of systems such as Shakey and SHRDLU, have their own drawbacks.

### Learning

The ability to learn is another key component of intelligence. By learning, we mean the process by which a system improves its performance on certain tasks on the basis of experience. We can view learning as a combination of inference and memory processes.<sup>22</sup> An agent with reasoning and learning capabilities is not limited by its designers' foresight, because it can tailor its behavior to new conditions as appropriate. Again, this enhances its adaptability. This capacity is crucial for things that cannot be known ahead of time, such as the layout of a user's home, or things that might change over time, such as the calibration of the robot's arm.

Learning is also useful for things that are simply difficult or tedious to directly program.<sup>17,23,24</sup> For instance, the most pervasive approach to industrial-robot task specification is still a combination of guiding and textual programming. A *teach pendant* (a handheld control device) deictically indicates the positions referenced in the program, which would otherwise be onerous to enter coordinate by coordinate. This is probably the simplest and oldest form of symbol grounding<sup>25</sup> in robotics.

Much robotics and control theory research has dealt with automatically modeling a mechanism's forward and inverse kinematics. A comparable amount of research has gone into elucidating the geometric transforms between sensor and motor subsystems. For example, Bartlett Mel shows how neural networks can progressively learn to move a robot arm to desired locations in a camera's visual field.<sup>26</sup> A lot of this work, most notably that of the CMAC (Cerebellar Model Articulation Controller) camp,<sup>27</sup> treats the problem as function approximation—for a given input, the system

should learn to predict the commonly observed output. This form of feedback is fairly rich in that it provides the exact desired response to the robot. However, even with this assistance, appropriately generalizing over the high-dimensional spaces typically used is difficult. This is especially true given the necessarily limited amount of training data—the robot has physical inertia, so each trial move takes a significant amount of real time.

Robotics researchers have also investigated a weaker form of feedback. In *reinforcement learning*, a reward signal tells the robot when it has done the right thing. One problem with this approach is that the robot does not really know how close it has come to the optimal solution or how much closer it might come if it varied some of its parameters more. Also, it is difficult to propagate such a reward signal across a series of actions to determine which ones truly produced the desirable outcome. Nevertheless, researchers have used reinforcement learning both to craft the transfer functions of individual behaviors<sup>24,28</sup> and to design the arbitration and sequencing logic for a given fixed set of behaviors.<sup>29</sup> Generally, the smaller the learning problem and the more frequent the reward signal, the faster such robots can learn.

Another major area of robot learning research involves automatically building maps of the environment. Typically, these systems are unsupervised and use various clustering techniques. Sebastian Thrun has refined and extended Hans Moravec's early work on sonar occupancy grids<sup>30</sup> into an elaborate probabilistic map construction and position estimation system.<sup>31</sup> Geometric models based on stereo vision have also been popular and have achieved some maturity.<sup>32,33</sup> Both approaches typically presuppose a weak model of the environment (walls and corridors), then use this bias to infer structural features from constellations of sensor readings. The only sort of feedback they use is a measure of consistency across different viewpoints.

Finally, learning can team up with high-level reasoning to further enhance an agent's competence. First, learning how to accomplish substeps is far easier than directly inducing a complete procedure. Second, a reasoning system can leverage even a relatively small fragment of learned knowledge by using it in conjunction with a wide range of other preexisting plan steps. High-level (symbolic) learning in general has been much investigated. Powerful algorithms exist

for empirically inducing representations from collections of examples. Also, deductive generalization techniques can use an underlying domain theory to explain a concrete case and, from this, derive a general solution for a whole class of problems. Combining inductive, explanation-based, and case-based approaches can significantly enhance the flexibility and adaptability of robots.<sup>17,34</sup> Unfortunately, owing to the difficulty of lower-level issues, researchers have left symbolic learning in robotics largely untouched until now.

### Communication

Linguistic communication is one of the main characteristics that distinguish humans from animals. To aspire to sentience—a fully human level of intelligence—an agent must be able to explain its reasoning and motivations. That is, its internal beliefs, desires, and intentions must be accessible to some extent by outside agents.

Accessibility lets an outside agent more easily specify commands or provide information to the robot. The agent can issue commands by either specifying top-level goals or directly enumerating whole sequences of requested steps (like an old command line interface). This is probably language's most significant direct use in robotics.

Yet language is also invaluable for grounding symbols. For instance, a particular room's name is not something the manufacturer can necessarily hard-code into a robot direct from the factory. But you could easily stand in a room and declare "This is the kitchen" to let the robot know what you mean by the command "Go to the kitchen." In this way, symbol grounding gives the human and robot a common vocabulary for interpreting commands and discussing alternatives.<sup>35</sup> You can also ground other things, such as adjectives ("red") or verbs ("follow"). You can even assign shorthand invocation designations, much like procedure names, to whole sequences of steps (for example,  $\text{find}(x) + \text{pick-up}(x) + \text{return-home} + \text{drop} = \text{clean-up}(x)$ ).

Of course, not all learning must be rote. You can also use language to specify an item's generic class or explicitly call out its important properties, to help the robot induce a recognizer for it.<sup>36</sup> For instance, "This is a bottle because of its shape" implies that the object's color, an easily observed overt property, is not relevant.

Perhaps most important, you can use lan-

guage to probe the robot's goals and plans. This can be invaluable for debugging, much as the dependency-directed backtrace of rule calls in an expert system helps justify its conclusions. With such a technique, the user can correct any underlying false misconceptions the robot might have. It also lets the user determine when the robot's planning system has hit an impasse and hence what additional knowledge and inference rules the robot needs. These uses are much closer to the psychological rationale for language we mentioned at this section's start.

### Integration & the articles in this issue

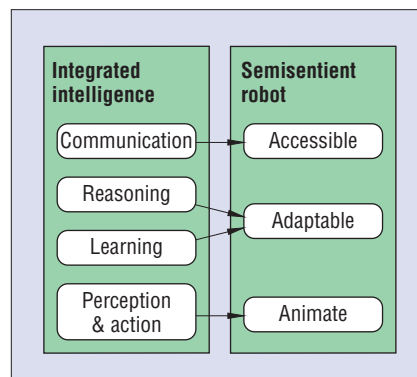
While no robot in this issue claims to be fully sentient, progress is occurring on various fronts. And, although the integration of action, reasoning, learning, and communication is complex, some consensus seems to be emerging. For instance, to allow reasoning but retain reactivity, layered control architectures are popular. Moreover, at some level (but possibly below the emotional and hormonal responses), reasoning is centralized to promote the robot's unity of purpose. Yet these researchers do not implicitly trust (the results of) reasoning. Their systems typically incorporate a lot of machinery to check the predictions made at regular intervals. They also provide facilities for debugging and partial replanning when the robot encounters unexpected events.

Furthermore, not all decisions result from general-purpose reasoning. In several systems, complex, special-purpose engines handle navigation, image processing, and speech recognition. Similarly, researchers tend not to use a single monolithic learning agency but rather sprinkle various types of learning throughout the robot's subsystems. Finally, reasoning and language appear to have a tight, almost Whorfian, coupling. Reasoning is most easily cast as a symbol-based system; language helps ground out these symbols, whether they represent objects or actions.

Such symbol grounding for natural language is the focus of Luc Steels' article. His Talking Heads project attempts to build a common vocabulary based on agents with a similar perceptual organization. It uses a shared attentional focus derived from the distinctiveness of presented objects. The projects with Sony Aibos and SDR humanoids attempt to extend this methodology to robots with more sophisticated sensor systems (including speech recognition). Steels also

suggests using the correspondingly more complicated motor systems of these robots (with many degrees of freedom) as a substrate for learning action labels.

The articles by Francois Michaud, Michael Beetz, and Hideki Asoh and their colleagues describe the state of the art in integrated robotic systems, all using variations on the hybrid deliberative-reactive control architecture. Michaud and his colleagues designed their Lolitta Hall robot specifically to address the AAAI Mobile Robot Challenge: register for the conference, schmooze, then give a talk. They use an intriguingly complex system of finite-state machines, reactive behaviors, and emotional modifiers



**Figure 1. The relation between integrated robot intelligence and the definition of semisentience.**

to schedule the robot's activities.

Beetz's Rhino also performs indoor navigation. Rhino has a sophisticated probabilistic sonar system for mapping and localization and employs a text-based language interface for accepting user commands. The robot uses an interesting clustering-based learning algorithm to learn common path segments, which the robot's reasoning system can then incorporate as substeps of an overall trajectory to enhance its adaptability.

Asoh's Jijo-2 similarly navigates through an office environment but uses full speech recognition to let a human specify travel destinations and to obtain relevant information from the human. Jijo-2's interaction capabilities include turning to the user on the basis of the direction of sound and detecting and recognizing the user's face. In addition, during supervised map learning, a human can verbally indicate relevant landmarks as the robot encounters them.

Ian Horswill examines how to bind linguistic entities to sensor data and what archi-

tectures support this naturally. His use of visual routines and markers is inspiring in its effectiveness and elegance. He also shows how a largely reactive architecture can perform pseudosymbolic reasoning and why such inference is useful for a robot. Although this research is very basic, his group has succeeded in building several robots that competently operate using these principles.

Finally, Stanislaw Lauria and his colleagues demonstrate in detail how a robot can interpret linguistic travel directions. This includes understanding what pieces of information are needed, requesting them when they are not apparent or ambiguous, and supplying default values as appropriate. The authors also derive an interesting set of motion primitives from a corpus of human interactions and suggest how to combine these primitives to yield higher-level named routines (such as "goto-library").

**T**o achieve a modicum of sentience, a robot must be animate, its behavior adaptable, and its motivations accessible. Figure 1 depicts the relation of these desiderata to the basic abilities of communication, action, reasoning, and learning. Creating such a semisentient robot touches on many of the central issues of AI: planning, natural language, vision, and learning. Although classical deliberative systems are eminently suited to natural languages, they have proved difficult to apply to real-time systems. Behavior-based approaches, on the other hand, are very responsive to their environment yet have problems dealing with symbols. Past attempts to integrate the two traditions have met limited success—the flexibility to accommodate and reason about new situations seems missing.

The merger of these abilities is important, because to achieve a true measure of sentience a robot must not only be alert and proficient at some real task, it must also provide a friendly, high-level interface to humans. Such an interface serves not only for telling the robot what to do but also for explaining how to do it and how to ground relevant symbolic terms in the robot's sensor and motor primitives. The articles in this issue explore new ways of combining communication with reactivity, reasoning, and learning. As will become apparent, although the resulting robots might not be fully sentient by human standards, they certainly provide enthralling demos—and perhaps a glimpse of what's to come. ■

References

1. D. Stork, ed., *HAL's Legacy: 2001's Computer as Dream and Reality*, MIT Press, Cambridge, Mass., 1997.
2. M. Brady, "Artificial Intelligence and Robotics," *Artificial Intelligence*, vol. 26, no. 1, Apr. 1985, pp. 79–121.
3. G. Walter, *The Living Brain*, W.W. Norton, New York, 1953.
4. H. Moravec, *Robot: Mere Machine to Transcendent Mind*, Oxford Univ. Press, Oxford, UK, 1998.
5. T. Winograd, *Understanding Natural Language*, Academic Press, San Diego, Calif., 1972.
6. P. Agre and D. Chapman, "Pengi: An Implementation of a Theory of Action," *Proc. 6th Nat'l Conf. Artificial Intelligence (AAAI 87)*, AAAI Press, Menlo Park, Calif., 1987, pp. 268–272.
7. R. Arkin, *Behavior-Based Robotics*, MIT Press, Cambridge, Mass., 1998.
8. R. Brooks, *Cambrian Intelligence*, MIT Press, Cambridge, Mass., 1999.
9. J. Connell, *Minimalist Mobile Robotics*, Academic Press, San Diego, Calif., 1990.
10. M. Mataric, "Environment Learning Using a Distributed Representation," *Proc. 1990 IEEE Int'l Conf. Robotics and Automation (ICRA 90)*, IEEE CS Press, Los Alamitos, Calif., 1990, pp. 402–406.
11. P. Maes, "The Dynamics of Action Selection," *Proc. 11th Int'l Joint Conf. Artificial Intelligence (IJCAI 89)*, vol. 2, Morgan Kaufmann, San Francisco, 1989, pp. 991–998.
12. J. Connell, "SSS: A Hybrid Architecture Applied to Robot Navigation," *Proc. 1992 IEEE Int'l Conf. Robotics and Automation (ICRA 92)*, IEEE CS Press, Los Alamitos, Calif., 1992, pp. 2719–2724.
13. E. Gat, "Integrating Planning and Reacting in a Heterogeneous Asynchronous Architecture for Controlling Real-World Mobile Robots," *Proc. 10th Nat'l Conf. Artificial Intelligence (AAAI 92)*, AAAI Press, Menlo Park, Calif., 1992, pp. 809–815.
14. J. Firby et al., "Programming CHIP for the IJCAI-95 Robot Competition," *AI Magazine*, vol. 17, no. 1, Spring 1996, pp. 71–81.
15. D. Isla et al., "A Layered Brain Architecture for Synthetic Creatures," *Proc. 17th Int'l Joint Conf. Artificial Intelligence (IJCAI01)*, Morgan Kaufmann, San Francisco, 2001, pp. 1051–1058.
16. K. Perlin and A. Goldberg, "Improv: A System for Scripting Interactive Actors in Virtual Worlds," *Proc. 23rd Ann. Conf. Computer Graphics (SIGGRAPH 96)*, ACM Press, New York, 1996, pp. 205–216.
17. L. Seabra Lopes, *Robot Learning at the Task Level: A Study in the Assembly Domain*, PhD thesis, Faculdade de Ciências e Tecnologia, Univ. Nova de Lisboa, Lisbon, Portugal, 1997.
18. G.R. Meijer, *Autonomous Shopfloor Systems: A Study into Exception Handling for Robot Control*, PhD thesis, Dept. of Computer Systems, Univ. van Amsterdam, Amsterdam, 1991.
19. S.J. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice-Hall, Upper Saddle River, N.J., 1995.
20. H. Levesque et al., "Golog: A Logic Programming Language for Dynamic Domains," *J. Logic Programming*, vol. 31, nos. 1–3, 1 Apr. 1997, pp. 59–83.
21. J.L. Kolodner, *Case-Based Reasoning*, Morgan Kaufmann, San Francisco, 1993, pp. 331–363.
22. R.S. Michalski, "Inferential Theory of Learning: Developing Foundations for Multistrategy Learning," *Machine Learning: A Multistrategy Approach*, vol. 4, R.S. Michalski and G. Tecuci, eds., Morgan Kaufmann, San Francisco, 1994.
23. J. Connell and S. Mahadevan, eds., *Robot Learning*, Kluwer Academic, Dordrecht, Netherlands, 1993.
24. K. Morik, M. Kaiser, and V. Klingspor, eds., *Making Robots Smart: Behavioral Learning Combines Sensing and Action*, Kluwer Academic, Dordrecht, Netherlands, 1999.
25. S. Harnad, "The Symbol Grounding Problem," *Physica D*, vol. 42, 1990, pp. 335–346.
26. B. Mel, *Connectionist Robot Motion Planning*, Academic Press, San Diego, Calif., 1990.
27. J. Albus, *Brains, Behavior, and Robotics*, McGraw-Hill, New York, 1981.
28. S. Mahadevan and J. Connell, "Automatic Programming of Behavior-Based Robots Using Reinforcement Learning," *Artificial Intelligence*, vol. 55, no. 2, June 1992, pp. 311–365.
29. P. Maes and R. Brooks, "Learning to Coordinate Behaviors," *Proc. 8th Nat'l Conf. Artificial Intelligence (AAAI 90)*, AAAI Press, Menlo Park, Calif., 1990, 796–802.
30. H. Moravec and A. Elfes, "High Resolution Maps from Wide Angle Sonar," *Proc. 1985 IEEE Conf. Robotics and Automation*, IEEE CS Press, Los Alamitos, Calif., 1985, pp. 116–121.
31. S. Thrun, "Probabilistic Algorithms in Robotics," *AI Magazine*, vol. 21, no. 4, Winter 2000, pp. 93–109.
32. N. Ayache and O. Faugeras, "Maintaining Representations of the Environment of a Mobile Robot," *IEEE Trans. Robotics and Automation*, vol. 5, no. 6, Dec. 1989, pp. 804–819.
33. L. Iocchi, K. Konolige, and M. Bajracharya, "Visually Realistic Mapping of a Planar Environment with Stereo," *Proc. 2000 Int'l Symp. Experimental Robotics (ISER 2000)*, 2000.
34. L. Seabra Lopes, "Failure Recovery Planning in Assembly Based on Acquired Experience: Learning by Analogy," *Proc. IEEE Int'l Symp. Assembly and Task Planning (ISATP 99)*, IEEE Press, Piscataway, N.J., 1999, pp. 294–300.
35. L. Seabra Lopes and A.J.S. Teixeira, "Human-Robot Interaction through Spoken Language Dialogue," *Proc. IEEE/RSJ Int'l Conf. Intelligent Robots and Systems (IROS 00)*, IEEE Press, Piscataway, N.J., 2000, pp. 528–534.
36. J. Connell, "Beer on the Brain," *My Dinner with R2D2: Natural Dialogues with Practical Robotic Devices* (working notes from AAAI 2000 Spring Symp.), AAAI Press, Menlo Park, Calif., 2000, pp. 25–26.

The Authors



**Luis Seabra Lopes** is a Professor Auxiliar at Universidade de Aveiro. He also leads the Intelligent Robotics Transverse Activity at the university's Instituto de Engenharia Electrónica e Telemática.

His interests include robot learning at the task level, spoken-language human-robot interaction, and service-robotics applications. He holds a licenciatura degree in computer science and a PhD in electrical engineering from Universidade Nova de Lisboa, Lisbon. He is a member of the IEEE. Contact him at the Dept. de Electrónica e Telecomunicações, Univ. de Aveiro, P-3810 Aveiro, Portugal; lsl@det.ua.pt.



**Jonathan H. Connell** is a research scientist at IBM's T.J. Watson Research Center, where he has worked on mobile-robot navigation, machine learning, wearable computers, natural language understanding, biometric identification, and vegetable recognition. He is a member of the Exploratory Computer Vision Group. He received his PhD in artificial intelligence from MIT. He is a member of the IEEE and AAAI and was an associate editor for *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Contact him at the IBM T.J. Watson Research Center, 30 Saw Mill River Rd., Hawthorne, NY 10532; jconnell@us.ibm.com.