Intelligent Mobile Robots: Perception of Performance

David P. Miller
MITRE Corp.
7525 Colshire Drive
Hayes Bldg
McLean, VA 22102
USA
dmiller@starbase.mitre.org

November 8, 1998

Abstract

How do intelligent robots interact with their environment? This paper will show that in some cases, it is the interaction with the environment that both defines and creates the intelligence. It is the synergy between the robot and the environment for which the robot has been *designed* that creates the intelligence. Environments that are outside the robot's area of competence are self defining as areas where the robot does not act intelligently. This is very similar with models on common-sense reasoning. Experimental evidence will be given that shows that intelligent behavior is not necessarily a product of deep reasoning about the environment.

1 Acting Intelligent

An intelligent agent is, by definition, an agent that does intelligent things. We usually define the accomplishment of some intelligent act as having the agent *intend* to accomplish the act, and then actually accomplish the act by some method that appears to be directed action. But what an agent should, or wants to, accomplish depends on the environment that the agent finds itself, and its long term goals. The agent will try to modify the state of the environment, or its state in the environment in service of these long term goals.

For example, a graduate student will eat whenever encountering food. A graduate student might also give the appearance of following a course of action to get an education by attending various seminars. However, in most cases, the graduate student is reacting to the notification that food will be served at the seminar. The attention given to the speaker is so that the student can best estimate when the talk will conclude and therefore be able to get in line early for the refreshments.

Seemingly intelligent agents often take a course of action that in a different environment, would not seem at all intelligent. When the agent is taken out of the environment in which they exhibit intelligent behavior, and are placed in a different context, they are often exposed to be somewhat less than truly intelligent. If it is a computer program that completely falls apart outside of its narrow domain then people often remark that the program is "Eliza-like", if it is a person who is incapable of operating outside of a narrow domain, then they are referred to as "idiot-savant". However, the difference between an Eliza-like program and a fully functional AI may not really be a difference of kind, but of extent.

A Montana farmer suddenly dropped onto the floor of the New York stock exchange is unlikely to perform intelligently by the standards of any professional trader. A nineteenth century stagecoach driver is equally unlikely to perform well when placed into the cockpit of a 747 on final approach, or at an international summit on arms control, or on the floor of a semi-conductor production facility. In fact, the overall performance of the driver would probably not be appreciably better than that of a chimpanzee or a potted palm. The only significant difference between the performance of our hypothetical anachronism and that of the potted plant would be that the nineteenth century driver might ask "how" or "what" or "help." The other difference is that give enough time and help, a year later the driver's performance should have improved significantly (as would the chimpanzee's) but the palm's performance would probably remain relatively unchanged.

So what does all this have to do with intelligent robots?

An robot's "apparent intelligence" is totally dependent on the environment in which the robot's performance is measured. As has been pointed out in in numerous studies, so-called general intelligence tests are culturally and economically biased. Beyond that, they are measuring certain qualities of how humans function in modern general society. They are no more relevant for measuring the capability to act intelligent while flying a 747 or operating on the floor of the stock exchange than is measuring the Ph of the soil surrounding our potted palm. Intelligence is reevaluated by an observer continuously. An observer will call a robot intelligent if the observer has not recently seen the robot do something obviously *unintelligent*.

2 Intelligent Robots & Robot Intelligence Tests

Most researchers would agree that very few robot's so far created act particularly intelligently. Given the discussion above, we could argue that it isn't the robot's that are necessarily the problem. The secret to making an intelligent robot is to make a capable robot, and then putting that robot in an environment where those capabilities are the definition of local intelligence.

Perhaps the most glaring example of the environment defining the intelligence of a robot occurred at the AAAI-92 mobile robot contest [Dean93]. At this event a dozen different robots from a variety of institutions compared timed performances at a couple different standardized tests. A large flat arena (approximately 30m in diameter) was walled off. Inside the arena a few dozen obstacles (large cardboard boxes) were randomly placed. Also inside the arena were ten 3m tall poles each 8cm in diameter. These were also randomly distributed. The first test was for each robot to be placed into the arena (one at a time), and with no a-priori knowledge try and visit each of the ten poles. A visit was defined as being within two robot diameters and then having the robot identify that it was at a pole, and which pole it was at. The poles could be individually marked by the robot-handlers in almost any way that they wished. The method most teams used was to use some sort of bar code on the pole.

The second test was to go visit three of the poles (selected by the judges) in order (determined by the judges). The robots could use any information they gathered during their runs in the first test to aid in the second round. However, this was not required.

Most of the robots used traditional AI systems for achieving this task. Most robots followed some search pattern that would guarantee covering most of the arena. Typically, they would scan the area near the robot for pole-like objects. These were determined by a vision system looking for tall skinny objects, or a sonar with a similar filter. When a suspected pole was located, the robot would approach the pole, verify that it was a pole, weave around any obstacles between it and the pole, identify the pole, and then announce it had found a pole [Congdon93]. Then the robot would continue on with its search pattern. Most robots also kept and updated some representation that described where each pole was found either topologically, geometrically, or in move space.

Only two of the robots were able to find all ten poles in the allotted time (twenty minutes). Three more robots found seven or more poles in the allotted time. Four out of the five of these robots used algorithms similar to that described above. The one exception was *Scarecrow* [Miller92b] which used a radically different approach.

Scarecrow could not execute the algorithm above because of its limited computation system. The robot had an electro-mechanical four bit processor made out of springs and relays. The system had its search algorithm hard-wired, had four bytes of ROM. The technique it used for the first test was a pseudo-random walk, always moving forward.

The robot had curved bumpers all along its front half. When contact was made it would turn left if it hit on the right side of the robot and turn right if contact was on the left. The bumpers were shaped to reflect the robot's turning radius so that the bumper would stay in contact with the obstacle until the robot had turned sufficiently to clear the obstacle. The top of the robot contained a conductive bar code reader. When the robot ran into an obstacle the bar code reader would also contact the obstacle. If the obstacle was a pole, then the reader would contact a metallic bar-code on the pole. This caused the robot to announce that it was at a pole and readout the ID # of the pole. The coverage of the arena was performed by random walk.

Despite the fact that Scarecrow actually had to come in contact with the poles in order to recognize them, and despite the fact that the coverage was basically random, Scarecrow found seven poles in the allotted time. As might be expected, it found the first five in the first few minutes and then spent the rest of the time picking up the next two. This performance put it in the top half of the contestants.

In the second test, the robots had to visit three selected poles in the order specified by the judges. A visit was defined as being near the pole in question and the robot signaling that it was at the desired pole. Scarecrow duplicated it random walk of the arena that it had performed in the first test. But when it ran into a pole it did a logical AND of the pole's ID with the ID in the top of its list in ROM. If there was a match, the robot signaled it had visited the desired pole and incremented the pointer into ROM. If there was no match, then the robot just continued on its path.

Scarecrow had the third best time of any robot on this test. While there was unquestionably some luck involved in the quality of the robot's performance, it none the less brings up some interesting questions on the incremental performance benefit of increased computation. These questions become even more significant when the audience's reaction is taken into account. When it was explained to people that the robot had no transistors, no memory and almost no computing, the question asked was "How does it store its map?" Dozens of people asked how it calculated how long to turn when it ran into something, and dozens more wanted to know how it homed in on poles.

The people asking these questions were AI and robotics researchers attending the AAAI conference. The robot had all of its pieces and parts fully exposed, and explanations were given in advance on how it worked. None the less, its performance was comparable to that of robots where a great amount of work had gone into making the system intelligent, and therefore Scarecrow, by analogy, had to be intelligent as well.

One researcher, who went over in detail how Scarecrow worked was heard muttering "Its random, but it worked as well as the planners, but it was random, but it worked..." Some people were very upset by the quality of the robot's performance. Any intelligence credited to this robot had to be due entirely to its performance. There was nothing in its structure to encourage such claims.

3 Intelligence, Performance, and Reactivity

Scarecrow is an extreme example of what has been happening in the field of intelligent robotics for the past half dozen years. Starting in the late eighties a series of relatively low-computation robots have been created. Their main difference from earlier systems is the reduction, or total lack, of internal state and world models. These systems have been programmed in a variety of special languages (e.g., Subsumption [Brooks86], ALFA [Gat91]) and are called many different things, but fall under the general rubric of reactive behavior control.

Behavior control robots usually have collections of tight feedback loops that closely couple sensors with actuators. There are few or no large deliberative processes in these systems, and as a result the robots usually react quickly to changes in their environment. While the actual number of behaviors is usually fairly small, the combinations of behaviors, combined with the fact that the active behaviors are constantly changing as the stimulus through the sensors changes, makes the observed behavior of these robots appear quite complex.

Since the required computation systems for the behavior control robots is small and inexpensive, the resulting robots are also often small and inexpensive. Practitioners of this discipline tend to produce lots of different robots that exhibit different sets of behaviors [Brooks89]. This can be contrasted with those researchers who work on more deliberative systems. Those robots are larger, and more expensive. Rather than coming out with a new robot or two every year, practitioners of deliberative approaches tend to make incremental improvements to their system, but keep the same basic hardware from year to year. Finally, Those doing behavior control robots tend to have robots that react in real time. Those doing deliberative systems tend to be severely compute bound in all their processes and are often unable to have their robots react in real time to a changing world.

The result has been astounding. The popular press, industry, and most of the scientific community outside of the AI field has latched onto the work of those doing behavior control. While the enthusiasm is somewhat understandable (these robots, whatever their capabilities, are much closer to being usable by outsiders then are the very complex deliberative systems), it is interesting that these outside observers have also credited these behavior control systems with much more intelligence than the deliberative systems.

The Behavior control systems are considered more intelligent because these robots "wear their intelligence on their sleeve." These robots react immediately to changes and do something about the changes. They may not give much thought about how to react, but in most cases they do the reasonable thing. The deliberative systems tend to be slow, and slow is often equated with stupid. When faced with a quickly changing world the deliberative system will go through great machinations on deciding what to do next. The result can appear incorrect for two reasons. The first is that the system will decide to do the optimal thing, which is often not the obvious thing and will go against the observer's instinct of what the robot should have done. The second, and more common problem is that the robot will choose to do something that made good sense when it started its computation, but by the time it tries to execute that act, the act is out of date.

The exact details are unimportant, but what is important is that to those outside of AI, performance is much more important than cognition in evaluating intelligence. In AI there is growing disagreement about this issue. Those arguing that cognitive processes are a requirement for true AI and for intelligence are starting to sound like those who argue that AI is impossible. Searle [Searle80] has argued that AI cannot ever succeed because if it did then we would have a mechanistic model of humans and humans are special and cannot be defined mechanistically. AI researchers doing deliberative robotics make a similar argument. Behavior control cannot be be a correct model of intelligence because the processing is to simple, and by analogy, would imply that we are simple.

4 Conclusions

Much of what people do is simple. It does not take a rocket scientist to walk down the street, go through doors, pick up rocks, etc. It takes a rocket scientist to build a rocket, but not to fly one (the fact that we normally have rocket scientists flying rockets is a cultural anomaly). Approaching picking up rocks the same way we have a robot play chess or be creative is almost certainly a mistake. Having a simple robot do much of the job of an astronaut [Miller92] is should not be considered an insult to the astronaut, and does not imply a major shock in the natural order of things.

Intelligent control of mobile robots means appropriate use of appropriate technology. Deliberation should be used as needed, but it is certainly not needed everywhere. We have experimental evidence to this effect. Similarly, it is probably unwise to assume that we can continue to add behaviors to a system that will allow it to play chess, be creative, and build a rocket. Some overall organizing scheme will almost certainly be needed, as will some way for the system to maintain internal models and perform simulations. Deliberative computation has a role in robotics. But for now, we should remember that you have to be able to roll before you can build.

5 References

{Brooks86} Rodney A. Brooks, A Robust Layered Control System for a Mobile Robot, *IEEE Journal on Robotics and Automation, vol RA-2#1,* March 1986.

{Brooks89} R. Brooks & A. Flynn, Fast, Cheap, and out of Control: A Robot Invasion of the Solar System, *Journal of the British Interplanetary Society, vol 42*, #10, pp478-485, October 1989.

{Congdon93} C. Congdon, M. Huber, D. Kortenkamp, K. Konolige, K. Myers, A. Saffiotti, and E. Ruspin, CARMEL Versus Flakey: A Comparison of Two Winners, in *AI Magazine*, vol 14 #1, Robotics Competition Conference Supplement, 1993.

{Dean93} T. Dean & R.P. Bonasso, 1992 AAAI Robot Exhibition and Competetion, in *AI Magazine*, vol 14 #1, Robotics Competition Conference Supplement, 1993.

{Gat91} ALFA: A Language for Programming Robotic Control Systems, in *Proceedings of the IEEE Conference on Robotics and Automation*, May 1991.

{Miller92} D.P. Miller, R.S.Desai, E.Gat, J.Loch & R. Ivlev, Reactive Navigation through Rough Terrain: Experimental Results, in the *Proceedings of the Tenth National Conference on Artificial Intelligence*, pp823-828, SanJose CA, 1992.

{Miller92b} D.P. Miller & J.Q. Milstein, Scarecrow the Robot, in *AI Magazine*, vol 13 #2, Robotics Competetion Conference Supplement, 1992.

{Searle80} John R. Searle, "Minds, brains, and programs," *Brain and Behavioral Sciences* 3 (1980) 417–424.