The Long-Term Effects of Secondary Sensing

David P. Miller

■ To integrate robotics into society, it is first necessary to measure and analyze current societal responses to areas within robotics. This article is the second in a continuing series of reports on the societal effects of various aspects of robotics. In my previous article, I discussed the problems of sensor abuse and outlined a program of treatment. However, despite the wide dissemination of that article, there are still numerous empty beds at the Susan Calvin Clinic for the Prevention of Sensor Abuse. Sensor abuse continues unabated despite strong evidence that there is a better way. In this article, I explore the age-old question, Why does the robotics community look down on efficient sensing systems?

Science and technology march on. Although society might try to stop them, it can only delay the inevitable. From the invention of fire to the creation of video games, technology has had its naysayers, particularly those being displaced by the new technology. The job of easing new technology into popular use falls on the shoulders of social psychologists.

On the Role of Social Science in Society

Society's current technological fear comes from emerging forms of robotics. The group most threatened isn't the factory worker but, rather, members of certain esoteric branches of computer science. They are worried about a disturbing trend that shows that as robots become more and more capable, they are using less and less of the theoretical work designed to create intelligent agents. Their fear that these high-capability robots will soon be universally accepted is not unfounded. Just as we got the United States to stop worrying and accept weapons of mass destruction (including commercial television), our plucky group of social scientists is now working on the societal acceptance of modern robotics.

However, the computer scientists have nothing to fear. Many have a constructive place in this exciting technology. Others can be retrained for productive work. For those remaining, especially those still caught in the web of expensive primary sensing, they should take heart in the situation of railroad firemen, those spirited individuals who were trained to stoke coal in steam engines and who continue to ride the rails on today's diesel engines. These computer scientists can rest assured that they will be supported in the style to which they have become accustomed.

In this article, I discuss secondary sensing. Secondary sensing is shown to be a key methodology for producing practical robotics, yet it is generally frowned on or, at least, relegated to the dust bin of "it's just engineering." What is secondary sensing and why it is unjustly ignored are the subjects of this article.

Introduction to Sensor Abuse

Find a nice quiet room with a comfortable chair. Sit back, relax, close your eyes. Nice, isn't it? You are experiencing a sensor-free, natural high. For years, this was the state in which all robots continuously existed. They were satisfied and so were we. Then, in the late 1960s, several young, poorly groomed robotics researchers started wiring sensor after sensor into their robot systems.

Automated systems that for years had reliably put the caps on beer bottles or wrapped cellophane around cigarette packs were becoming tremendously more complicated. Cameras were added, along with large computers, special lighting, more computers, software engineers, more computers, and so on. Simple automation was turned into complicated, sensor-abusing robotics, which led to more complicated production and today's high prices on packaged beer and tobacco products. Such is the unfortunate price we all have to pay, but what is worse is that almost all the research that leads to serious sensor abuse is government sponsored.

What separates sensor abuse from traditional government waste is that the related professional organizations also support the program. Sure, there are rival gangs of sensor abusers. The *Libs*[1] believe that one or two sensors is more than adequate as long as your perception algorithm is depth first and nonterminating. The *Bloods*[2] don't think about what they are sensing, and they believe that their robots shouldn't think much about it either. However, despite their different philosophies, both the Libs and the Bloods believe that the government should continue to subsidize the tools of sensor abuse no matter what the cost to society.

In addition to propagating technology that supports sensor abuse, the robotic research community has teamed up with the intelligence community[3] to wage a covert campaign against secondary sensing. This situation is most unfortunate. In my opinion,[4] secondary sensing offers a low-cost, practical way of efficiently maneuvering a robot through a complex environment without falling into the unending cycle of government handout and sensor abuse.

Secondary Sensing

"Sensors aren't slow, algorithms are."

This popular, catchy phrase of the sensor-user community might technically be true, but it misses the point of the sensor-abuse problem. I have heard[5] that as much as 97 percent of a robot's processing power can be engaged in sensor processing. Removing all algorithms would leave the robot completely nonfunctional, but removing all the sensors would allow the robot to still accomplish many useful tasks. Fortunately, we don't have to go this far. With the careful application of second-hand sensing, we can have a robot maintain its full capability but actually increase its performance.

Most sensors are unnecessary. The Libs say that in-depth vision sensing is necessary for intelligent robots because people are intelligent and have vision sensors. This reasoning is faulty. Although people might be intelligent and have vision sensors, what makes them act intelligent is their perceptions, not the sensor. It is a fact[6] that people form strong perceptions with no data whatsoever. I do it all the time.

It is the conversion of sensor data to accurate perception that takes millions or billions of computer instructions to achieve. This fault is not with perception but with the complexity of most sensory data. The Bloods agree with this analysis but make a misguided response. Basically, they look for simple sensors (which is good), but then they make up in quantity what they reduced in complexity (which is bad).

What robotics needs is simple sensing that conveys complex ideas. Fortunately, we only have to take a quick look at how biological systems cope in nature to get a clear idea of where the solution lies…in secondary sensing.

What robotics needs is simple sensing that conveys complex ideas.

When a human is in a natural setting (behind the wheel of a Honda Accord, going 20 miles over the legal limit on a 2-lane road and trying to unwrap the imprecisely laid cellophane from a pack of cigarettes and simultaneously get the overtightened bottle cap off the next beer), neither nature nor the Federal Highway Commission (FHC) would expect him or her to be able to interpret the complex images and kinematic sensory output that would be needed to determine the layout of the road ahead. FHC has had a trained professional go down every road as it is laid at the speed for which it was laid and convert this huge amount of complex sensory data into a small set of pithy, easily parsable road signs that are all that we can be expected to deal with during high-speed travel. The driver does not have to be able to determine the rate of turn in the road or separate road from shoulder. FHC has already done it. FHC then masticates the sensory data into small, easily swallowed pills that can easily be digested by the computational capabilities of the most average of drivers. Bright-yellow reflective lines separate road from shoulder. Iconic and text representations tell the driver that the road turns left or right. The driver doesn't need to be able to see if the cross-traffic has a stop sign; he or she knows: He or she has been given the predigested sensing of the FHC technician in the form of an easily machine-parsable sign that says that the cross-traffic must stop. The driver doesn't need second sight; the driver has secondary sensing.

This is sensing the way nature intended it—no muss, no fuss, no strain on the brain. Secondary sensing doesn't rely on a particular sensor—even vision can sometimes be used in a secondary-sensing system. Secondary sensing minimizes the work in sensing by exploiting prior knowledge, offline processing, and implicit information in explicit environmental cues. What is so confusing about the sensor-abuse community is that although it is the first to sue the local government if a street is not properly outfitted with secondary-sensing cues, if anyone comes in with a robot that relies on secondary sensing—bar codes, beacons, GPS (global positioning system), or even written directions—it is referred to as cheating.

Clinical Experiments

More interesting than secondary sensing is the reaction of the average member of society to users of secondary sensing. To study this phenomenon, an experiment was conducted with a group of upstanding AI citizens who were placed in close proximity to a group of engineers evenly divided between primary sensor abusers and secondary-sensor users.

This is sensing the way nature intended it—no muss, no fuss, no strain on the brain

To exaggerate reactions (making them easier to detect), all the participants were locked in a large room for 72 consecutive hours. A barely audible but incessant 60-hertz hum was played in the room to help bring emotional reactions to the surface. Additionally, one of my assistants, Dr. Murphy Laws, quietly went around disconnecting wires, randomly typing on keyboards, and so on, to introduce an added level of stress to the tasks that the participants were trying to complete. Interactions among the participants were recorded and videotaped.

To get meaningful reactions, the experiment was conducted following the Shitake protocol.[7] The engineers were told that they were taking part in a performance contest. The AI observers were informed that they were attending a technical conference. The latter deception not only allowed us to get valid reactions but also helped to fund the experiment. Finally, we gave Dr. Laws a cloak of invisibility to help him covertly carry out his part of the experiment.

As part of the rules on doing potentially debilitating psychological experiments on humans, certain medical facilities had to be collocated with the experiment. Because my colleagues and I all have Ph.D.'s, no other doctors were required on site. However, some of the emergency equipment was quite bulky and difficult to hide. We also had a large amount of communications and recording equipment to deal with. Believing in the maxim of hide in plain sight, we simply taped a few random electronics parts to the wheelchair and other large equipment, stacked it all together, and just told people that we were entering the equipment in the robot contest. No one was the wiser.

Experimental Trials

Our experiments consisted of the environmental setup and two major trials. We explored two instances of secondary sensing. The first instance used a completely mechanical secondary-sensing system made up of human-placed positional beacons. The second instance used a human assisting a robot. In the latter run, the human condensed various sensor information into a binary response.

Experiment One

The first run had as its goals the quantitative establishment of the extent of the bias against blatant secondary sensing. A contest was set up where robots would have to escape from a closed office when one of the doors was suddenly opened. The doors were specially marked with visual bar codes so that a vision system could be used to identify when a door was open (because its sign would then disappear from view). There were no rules requiring a robot to use the visual cues. The robots would then try and make their way to a goal wall a dozen meters away. The robot that made it in the best time would win.

Thirteen robots and our pile of emergency medical and recording equipment were entered into the contest. At the end of the first round of timed runs, the five fastest robots completing the course were, in order of slowest to fastest, (1) XAVIER, which used wheel encoders, bump sensors, a ring of sonars, a laser-range sensor, and a color camera; (2) FLAKEY, which used wheel encoders, bump sensors, a ring of sonars, and a camera; (3) SCIMMER, which used wheel encoders, bump sensors, a ring of sonars, infrared proximity, and a laser range finder; (4) ALFRED, which used wheel encoders and a ring of sonars; and (5) ISAAC, which used forward sonars and GPS beacons.

The difference between the color camera on XAVIER and the GPS beacons is huge. The color camera yields a megabyte of information for the processor to chew on. The GPS beacons give the processor three numbers (position and orientation) that are not open to interpretation. Color vision is a primary sensing capability, but the usefulness of the beacons is derived completely from secondary sensing; someone else has gone through and positioned the beacons so that the numbers they generate are meaningful. No real interpretation is needed at the time that the sensors are actually used by the robot. The results from the beacon sensors can be fed directly to the control algorithm.

To some extent, wheel encoders, bump sensors, and sonar are secondary sensors. Their value is maximized when environmental assumptions say that anything they see is an obstacle, and anything that they don't see is free space. However, because of their inherent lack of robustness, there is a tendency to use them more in a primary sensing mode, where they are combined with spatial representations and statistical methods, often leading to drowsiness and uncontrollable twitching.

Not surprisingly, our subjects, who had spent the previous 72 hours sleepless, trying to overcome the effects of our Dr. Murphy, were not pleased with these results. After lengthy consultation among themselves, they decided to penalize ISAAC for relying on secondary sensing. Points were deducted, and ISAAC was ranked at the bottom of the robots that successfully completed the course (tenth out of the 14 that attempted the course). A finals round was then held for the top scoring (not necessarily fastest) contestants, but the results from this run are irrelevant to the subject of this article.

Experiment Two

The first experiment showed that our subjects were offended and frightened by overt secondary sensing. The question remained, How far could we go with covert secondary sensing? This second experiment was designed to find the answer to this question.

The second task given to the contestants was to wander through an office maze and locate a coffeepot. The robots were given a map and were told what quadrant the coffeepot was in, but they were not told where they were starting. The assignment was to find the coffeepot and deliver it to a goal location specified on the map.

A provided map is the epitome of secondary sensing. Someone, with all the time in the world, has precompiled aerial views of the world into terse, easy-to-parse information that is only a memory access away. What's through the next doorway? It's difficult to sense directly, but with secondary sensing and a map, you can easily get a good idea. However, somehow, AI has convinced itself that a map is not a sensor at all. It is an internal representation—on the same scale as having a knowledge of physical laws and memorizing the works of Nietsche. Because of this conviction, a map is considered, in AI terms, knowledge that somehow elevates it beyond the status of a database of secondarily sensed data.

Eight contestants attempted this task. None of them successfully completed the task. Two robots did succeed in making it to the coffeepot and go on to the goal location, but both required help in identifying the coffeepot.

ISAAC had the fastest run. It immediately located itself in the map using its GPS sensors. Its trainers then gave it a strong bit of secondary sensing by telling the robot the location of the coffeepot. ISAAC then planned a path and navigated its way directly to the pot. It then went directly to the delivery location.

ALFRED stored its map data as a topologic representation and as procedural information. When it started its run, it followed a wall, noticing left and right turns that it made as well as hallway intersections. This information was matched against its map data until a unique match could be made. This match identified the robot's location. The robot then started a search pattern to find the room with the coffeepot. During this phase of the contest, the robot relied on a secondary-sensing strategy. When it entered a room, it would wait for its human trainer to identify whether this room was the one with the coffeepot. If it wasn't, ALFRED would go on to the next room. If it was, ALFRED would then proceed directly to the delivery room.

ALFRED and ISAAC both used effective strategies for accomplishing this task. Both relied almost exclusively on secondary sensing. The preprocessed map was used by both. ALFRED used motor-command history patterns matched into its map to do localization. Tracking motor-command history is a well-established method of fending off sensor abuse. ISAAC used direct secondary sensing through its GPS beacons to do localization. This method was so fast and effective that in articles and papers written after the contest, ISAAC was not even considered to have accomplished self-localization. Both systems relied on a human to tell them what room the coffeepot was in. ISAAC was told at the start of its run. ALFRED was told during its run, when it

Not surprisingly, our subjects, who had spent the previous 72 hours sleepless, trying to overcome the effects of our Dr. Murphy, were not pleased with these results.

was in the correct room, making it appear that the person was a virtual vision sensor. However, there is no doubt that either robot could have functioned well if it was told the location of the coffeepot at any time during its respective run.

Many of our subjects were upset with both of these robots. However, because these robots were the only two to complete the task, there was not a whole lot they could do. Nevertheless, although willing to give some leniency to the subtle use of secondary sensing in ALFRED, the decision was made to penalize ISAAC for overt secondary sensing. ISAAC was put way down in the pack, moving up contestants that although didn't complete the task, at least did not use secondary sensing.

Conclusions

Secondary sensing is a set of methods that allow a system to efficiently exploit sensor and domain knowledge with maximum benefit and minimum computation. However, primary sensing, without concern for the features of the task and domain, is inefficient and usually ineffective. Secondary sensing is a key methodology for building practical systems in the future. It has been used effectively in current systems such as ALFRED and ISAAC to outperform other robots that have behind them more computers, more sensors, more complexity, higher cost, and more labor. However, many of our subjects believe that the end should be modified to justify the means.

These experiments on researchers' views of secondary sensing have been instructive. It is obvious from the results that most researchers are repulsed, on some deep fundamental level, by the idea that research should lead to practical application. Narrow research focus discourages system integration—secondary sensing is more of an integration issue than a sensor-development issue.[8] To keep practical application to a minimum and ensure the continuation of depth-first research practices, some researchers are willing to value a robot unsuccessfully, using vision[9] more highly than one successfully integrating beacons into a navigation strategy.

However, this attitude is not unique to robotics researchers. It permeates all levels of American society, from train firemen to the highest levels of government. If we cannot overcome our aversion to actually doing something practical, then we might as well turn ourselves over to the sensor pushers and other scum on the street. Those that cannot learn from robot contests are doomed to repeat them.

Notes

1. This term is short for deliberators. They are headquartered in Northern California but can be found almost anywhere.

2. This term is short for bloody reactivists. They are headquartered in the Northeast but have infiltrated most of middle-class society.

3. This combination makes up the intelligent robotics community.

4. It is my opinion that my opinion is as good as most authors' facts. Therefore, opinion and fact are used interchangeably throughout this article.

5. My opinions are based heavily on what I have heard. Therefore, hearsay and fact are used interchangeably throughout this article.

6. At least, it is a strong opinion.

7. The Shitake protocol is a specialized form of the mushroom protocol, wherein the experimental subjects are kept in the dark and fed a lot of BS. In the Shitake protocol, BS is supplemented with late-night servings of Mu-shi vegetables and steamed rice.

8. Oddly enough, although not able to integrate worth a damn, most serious sensor abusers can differentiate difficult things, such as df/dw_H, yielding $a^2 + b(b \; 1 \; l)2al\mathrm{cot}u/a^2 + (b + l)^2$.

9. Vision can successfully be used in a secondary-sensing role on a robot. For example, ALVIN uses knowledge of highway markings, a specially trained neural net for region following, and (shudder) a camera to do robust automated highway traversing even in the face of road kills.

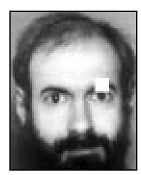

David P. Miller received his Ph.D. from Yale University in 1985. He has been part of the computer science faculty at Virginia Polytechnic Institute and State University and supervisor of the Robotic Intelligence Group at the Jet Propulsion Laboratory. After a two-year residency in the AI Lab at the Massachusetts Institute of Technology, Miller is now the director of the Behavior Modification Program at the Susan Calvin Clinic for the Prevention of Sensor Abuse. He also serves as a principal scientist for The MITRE Corporation in McLean, Virginia.