Design for an Autonomous Vacuum

Erann Gat
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109
gat@robotics.jpl.nasa.gov

This missive is being written for a workshop on autonomous vacuum cleaners. The problem statement reads as follows:

> ...this symposium will concentrate on AI as applied to a physically instantiated robot for vacuuming household floors. The target problem is to autonomously vacuum a living room while doing the right thing with furniture, trash, pets, etc.

I will argue informally that this is not an AI problem at all. None of the techniques currently being investigated by researchers can produce effective vacuuming behavior better than the simple expedient of robust mechanical design. In particular, I will argue that contact sensors which cover the entire body of the robot, together with a sealed drive mechanism and some means of self-localization, are both necessary and sufficient to produce robust autonomous vacuuming.

Let me begin with some assumptions.

Assumption 1: A robot vacuum cleaner will have to move. This may seem tautological, but it is not at all obvious that the most efficient solution to cleaning carpets is a mobile robot. It may be more economically viable, for example, to construct a carpet with a network of tubes built into it which carry away dirt continuously like a swimming pool filter. (Swimming pool cleaning, incidentally, makes an interesting case study for anyone interested in constructing autonomous vacuum cleaners, since autonomous mobile robots for cleaning pools are common. It should be noted, however, that cleaning a swimming pool is generally substantially easier than cleaning a carpet.)

Assumption 2: A robot vacuum cleaner will have wheels. The only alternative is legs, and all indications are that it will be some time before these can be made economically viable for consumer products.

Assumption 3: A robot vacuum cleaner will have to operate on bare floors and carpets of various colors, textures and thicknesses. This may not be true in all circumstances, but a robot which was restricted to a single surface would certainly be substantially less useful than one which met this requirement.

Assumption 4: Among the obstacles that the robot may encounter are non-somatic obstacles such as clothing or phone cords. There may also be darkly colored things that do not reflect infrared light or laser light very well. There may also be soft things that absorb ultrasound radiation. There may also be shiny things which reflect ultrasound in a specular manner. There may also be obstacles which are colored very similarly to the carpet. These were not specifically listed in the problem statement, but it is not unreasonable to suppose that such things will be encountered from time to time.

Assumption 5: The robot is not expected to move furniture and other object, but rather to vacuum around them. (On the other hand, some objects, like pets, may move about on their own.)

I will argue informally that under these assumptions it is necessary for a robot to be equipped with contact sensors over its entire body if it is to operate reliably. Furthermore, I will argue that these sensors, together with the ability to self-localize, are sufficient to produce reliable vacuuming behavior.

Let us begin by considering the usual solution to the problem of collision-free motion. A (wheeled) robot is first equipped with a number of sensors. These typically include such things as wheel encoders, sonars, infrared proximity sensors, inclinometers, and contact sensors. A well-endowed laboratory might add a laser range finder or a vision system. These six sensor modalities represent the vast majority of sensors currently in use on mobile robots in research laboratories.

Most of the commonly used sensors and strategies for avoiding obstacles will fail under the stated assumptions. None of the sensor modalities commonly in use can reliably detect all of the classes of obstacles which are commonly found in household environments. Sonars fail on soft objects and specular objects. Contact sensors fail on non-somatic obstacles. Emissive sensors like laser ranger finders fail on dark objects (and are expensive and hazardous besides).

Vision is conceptually the most promising obstacle sensing technology. We appear to have an existence proof of the adequacy of vision for vacuuming, but attempts to use vision to find obstacles so far have met with limited success. Horswill has suggested using vision not to find obstacles but to find carpeted areas. (Carpeted areas are not necessarily the complement of the areas occupied by obstacles.) However, his matched-filter approach tends to work only on one class of carpet at a time, and would probably have a very difficult time on patterned carpet.

Even if we had a vision system as good as a human's this would still not be enough. Humans rely not only on vision, but also on tactile and auditory feedback. It is not uncommon to vacuum up a paper clip, for example. Such minor failures of the human vision system are usually detected (in the case of vacuuming a paper clip, by hearing a loud rattling noise) and corrective action is taken. It is likewise not uncommon to have to vacuum in a dark corner or under the bed where you can't see; in such cases one typically has to push the vacuum around gently until it bumps up against something.

Wheeled robots have an additional constraint which they must deal with. Small, flexible objects like socks can get caught in the robot's drive mechanism and jam it. While this problem can be addressed by monitoring the drive motor current, it is important to distinguish this situation from the more benign condition of being in contact with a heavy object. In order to make this distinction, the robot must be able to tell when it is touching something anywhere on its body.

Of course, the most straightforward way of addressing this problem is to design a drive mechanism which is covered in such a way that foreign objects cannot enter it. It is impossible to seal a mobile robot's drive mechanism completely, since the wheels must be at least partially exposed in order to make contact with the ground. However, there is at least one commercial robot design (currently in prototype) which comes very close to the ideal. This design (from Real World Interface) consists of a cylindrical synchrodrive robot with a housing that covers the gear train for each drive wheel. A bottom plate covers the main drive mechanism. There are three circular hole in the bottom plate which expose the wheels. Each wheel is fitted with a circular cover plate which matches the contour of the hole. This mechanism is still vulnerable to jamming by very fine obstacles (a discarded piece of dental floss, for example) but should be robust in the presence of socks.

This robot has a second unique design feature. It can be equipped with contact sensors that cover the robot completely from top to bottom. The clearance between adjacent sensors is on the order of a millimeter. A strategically placed kitchen knife might be able to thwart this sensor, but that possibility can probably be safely discounted. Most obstacles that are small enough to penetrate the space between these contact sensors are not rigid enough to do any damage.

Such a design comes very close to the sort of reliability we require. Any obstacle it encounters which does not trigger the contact sensor can simply be pushed out of the way. A jammed drive mechanism can be detected by an increase in motor load with no forward motion and no detected obstacle. Such a robot, if equipped with some means of self-localization and a vacuuming attachment, can reliably and safely vacuum an area either by doing a random walk, or by employing any one a of a number of area-covering traversal algorithms.

The requirement of being able to self-localize may seem like a severe drawback, but it is not. There is at least one off-the-shelf beacon system which provides absolute position with high accuracy at high bandwidths. While this has historically been considered "cheating" in the AI community, the market typically has been insensitive to such ivory tower ideals.

The bottom line is that autonomous vacuuming is probably not nearly as interesting a problem as the AI/Robotics community would like to think. A well-designed set of bumpers and a beacon system will probably solve the problem effectively, while any other solution is likely to be unreliable.