

MASSACHUSETTS INSTITUTE OF TECHNOLOGY  
ARTIFICIAL INTELLIGENCE LABORATORY

Working Paper 302

November, 1987

## MIT Mobile Robots - What's Next?

Anita M. Flynn and Rodney A. Brooks

### Abstract

The MIT Mobile Robot Project began in January of 1985 with the objective of building machines that could operate autonomously and robustly in dynamically changing environments. We now have four working robots, each progressively more intelligent and sophisticated. All incorporate some rather novel ideas about how to build a control system that can adequately deal with complex environments. The project has also contributed some innovative and creative technical solutions in terms of putting together sensors, actuators, power supplies and processing power into whole systems that actually work. From our experiences over the past two and a half years, we have gained insight into the real issues and problems and what the goals should be for future robotics research. This paper gives our perspectives on mobile robotics: our objectives, experiences, mistakes and future plans.

---

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the research is provided in part by the University Research Initiative under Office of Naval Research contract N00014-86-K-0685, in part by an IBM Faculty Development Award, in part by a grant from the Systems Development Foundation, and in part by the Advanced Research Projects Agency under Office of Naval Research contract N00014-85-K-0124.

## 1. Introduction

### 1.1 Setting Goals For a Research Project

It is useful to state long range goals for any project because it provides both a way to explain current research and also helps to define what the next steps should be.

What are the goals of most mobile robot projects? Some are aimed at developing flexible automation technology, such as parts carts in cleanrooms and factories. Other projects are funded in the hopes of producing machines that will be useful in hazardous places such as nuclear power plants or battlefields.

At the MIT Artificial Intelligence Laboratory, the goals of the mobile robot project are greatly influenced by the goals of the AI Lab in general, which are twofold. One goal is to gain knowledge about the nature of intelligence, and implementing theories through computation is deemed the most appropriate way because it ensures intellectual honesty. The second goal is to make computers more useful to society, and creating more intelligent computers will make them more useful to people.

We claim that robotics is a superset of AI and that it will force AI to deal with the real world (this keeps us honest). We build our mobile robots in order to investigate intelligence and therefore learn about what it takes to build a system which we would consider clever. In addition, the second goal of our research is to make robots more useful to society. This can be achieved by making robots widely available to large segments of the population. Lowering costs for a given level of intelligence is one way to increase availability.

We have developed two strategies for achieving these goals. The first is to attempt to build *autonomous creatures*, in an effort to produce intelligent systems which operate in the real world so that we can learn more about the nature of intelligence. The second is to develop *mass producible, low-cost robots*, as robots can be more useful to people if more people have them, even if they have only limited intelligence. This idea follows from observing how microprocessors were first introduced into the public domain as cheap technology with very limited capabilities, but gradually spawned enough ideas and work on new applications that now they are definitely regarded as of great use to society.

The effect of our efforts along these directions has been the development of a series of mobile robots and some innovative ideas for future projects.

### 1.2 Autonomous Creatures

The goal of building *autonomous creatures* has led to the idea of building a control system, based upon task achieving behaviors, that can deal with both multiple goals and multiple sensors [Brooks 1986]. This control system has been implemented in various ways on four robots: Allen, Herbert, Tom and Jerry [Brooks 1987], [Connell 1987]. At present, they are all capable of autonomously wandering about and avoiding obstacles or following walls [Brooks and Connell 1986], with either sonar or infrared sensors. Herbert will eventually push beyond these levels of behavior and exhibit such complex actions as wandering the office

areas of the laboratory, finding soda cans on people's desks, retrieving them and returning them to a known location. The scenario set up for Herbert addresses four basic issues in mobile robotics research: recognition, manipulation, obstacle avoidance and navigation. Clearly, this line of work is attempting to build more clever robots.

### *1.3 Cheap, Mass-Produced Robots*

The goal of building *mass-producible, low-cost robots* has led to the concept of gnat robots [Flynn 1987], in which a complete robotic system is integrated onto a single piece of silicon. By using silicon micro-machining technology that is compatible with traditional integrated circuit processes, sensors, brains, actuators and power supplies can all be integrated on chip. The result of developing such a robot would be tremendous reductions in cost, due to the mass production possibilities of IC manufacturing. Additionally, vast new areas of application open up, because of the small size of such robots.

## **2. Past Experience**

### *2.1 A Different Approach*

We began our project by building Allen, which was designed to try out some novel ideas about robot control systems and planning models. The philosophy driving our project runs orthogonal to traditional implementations of mobile robots in that there is no central place in the robot's intelligence system which performs either supervisory control or sensor fusion into a global map. Rather, the control system is distributed and much of the intelligence is often found in the peripherals. This philosophy is described in more detail in [Brooks 1987b].

The absence of a control supervisor allows us to incrementally build layered control systems where each layer provides for increasingly sophisticated behaviors. This design for a brain was first reported in [Brooks 1986]. We call it the *subsumption architecture* because higher levels can subsume lower levels when triggered, but lower levels continue to work should higher levels fail or simply choose to remain inactive.

Starting from such very simple survival type behaviors, more complex behaviors are then gradually developed. What can be achieved is "insect level" intelligence and [Brooks 1986b] points out that evolution devoted much more time to these sorts of problems than to higher level reasoning. Each layer in the subsumption architecture is composed of a set of finite state machines which send messages to each other. We call such finite state machines *modules*. Designing the control system consists of "wiring up a brain" by assigning inputs and outputs between various modules.

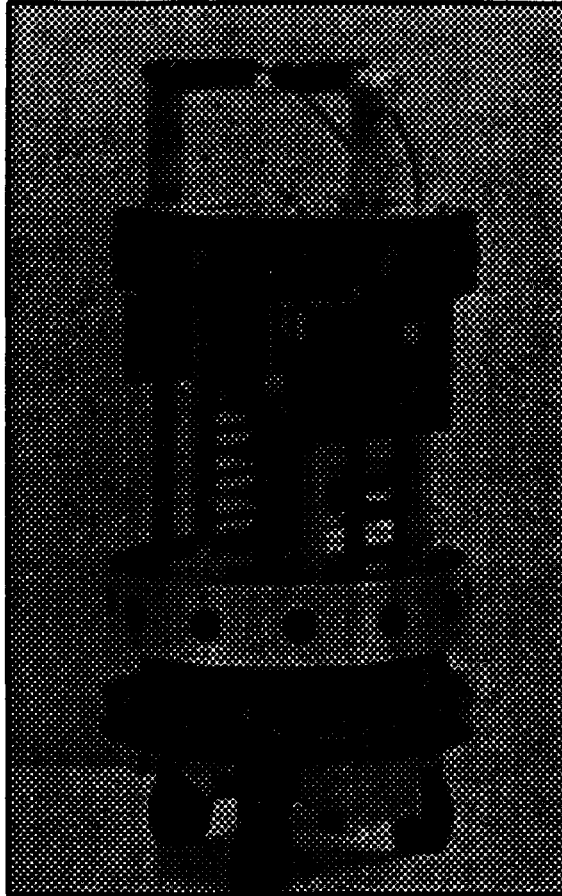


Figure 1. Allen used 12 sonar sensors for obstacle avoidance. Data was transmitted over a cable to an off-board Lisp machine running a subsumption architecture simulation. Allen is still used as a testbed for ongoing experiments in new sensor and control algorithms.

---

## 2.2 Allen

Allen, shown in figure 1, was built to try out some of these ideas. A three layered subsumption architecture was implemented on a Lisp machine. Via a cable, the Lisp machine collected sensor data from the robot and issued appropriate actuator commands. Allen had a ring of 12 sonar sensors mounted on a synchronous drive base which was propelled by two motors, a drive motor and a steering motor. All wheels turned in unison. The lowest level of the subsumption architecture created an *avoid* behavior in which the robot would sit in the center of a room and runaway if someone approached it. The *wander* behavior was built on top of *avoid* and forced the robot to wander randomly but in addition, enabled it to runaway if either someone approached it or it found itself wandering into obstacles.

The third level exhibited an *explore* behavior in which the robot would pick a destination and then try to go there, circumnavigating a person should someone step in front of it. We found the subsumption architecture to be very flexible and a variety of behaviors could be substituted quite easily, such as *door finding*, etc. These extensions were reported in [Brooks and Connell 1986].

Noticing that the finite state machines comprising these particular control systems were in fact, quite simple, we paused before adding more levels of behavior to Allen, and went ahead with two other projects which implemented this subsumption architecture directly in hardware in two separate forms. Both projects have enabled us to build fully autonomous robots, free from the cable constraint that fed Allen.

### 3. Current Progress

#### 3.1 Herbert

The Herbert project was begun to illustrate that a Lisp machine implementation of the subsumption architecture was essentially overkill for the computation that was actually being performed on Allen. Instead, Herbert's brain is composed of a network of 8-bit microprocessors that physically realize the composition of modules that previously had been written in software on the Lisp machine. The plan for the Herbert project was initially outlined in [Brooks, Connell and Flynn 1986] and more recent progress is discussed in [Brooks, Connell and Ning 1987].

Herbert exhibits roughly the same levels of behavior as Allen at the present time, but his microprocessor network emphasizes the fact that there really is no common memory, bus, switch or central controller. Rather, the processors are wired together in the form of a patch panel with outputs of some modules feeding inputs of others, and with outputs of modules in higher levels at times reaching down and suppressing inputs to modules in lower levels. Having this network of small processors onboard allows Herbert to run autonomously whenever he is powered up.

Herbert uses infrared proximity sensors instead of sonar sensors for obstacle avoidance as this gives him faster reflexes. Although at this stage Herbert exhibits only two fundamental behaviors, *avoid* and *follow walls*, he will soon be able to perform more interesting actions as he will have an arm incorporated into the control system shortly.

The goal in life chosen for Herbert is that he will wander around and collect soda cans and take them back to a central depository. Although situated in a light scenario, Herbert's existence in the context of collecting soda cans explicitly targets four major areas in mobile robotics research. First, the robot must be able to find the cans, therefore requiring that he solve a recognition problem. Second, the robot must be able to maneuver his way to the target without injuring himself in the process. This means he has to have some obstacle avoidance skills. Third, he has to be able to grasp the can, requiring manipulator control and end effector sensors. Fourth, he has to know how to get back to the soda can depository and this necessitates a capability for path recognition and navigation.



Figure 2. Herbert, the Collection Machine, has an onboard parallel processor, a lightweight arm, infrared proximity sensors, a laser light striper and an infrared hand sensor, all designed so as to provide him the capability of recognizing and collecting soda cans.

---

In keeping with the driving philosophy of no central control, all of these actions will be performed in a distributed fashion. A recently developed light striper vision system will be able to provide depth maps in real time to enable parallel recognition of a number of different types of objects. Each object class has its own special processor to recognize it. When the table-like-object recognizer fires, a table-approaching behavior activates. Being close to a table brings any soda cans thereon into the rough size range expected of soda cans. If a can is detected, the robot centers itself on the suspected can. If proximity sensors detect obstacles during this process they trigger avoidance behaviors. Safe arrival to the destination triggers the arm to scan along the tabletop or floor. Ranging sensors mounted on the hand fire when they detect a raised object of the right size. This triggers a grasping action. When the arm senses an external force on the hand, the behavior for navigating back to the depository is initiated. Note that nowhere in this scenario is it necessary to have a central controller. The arm does not know where in the room the soda can lies



Figure 3. Here the subsumption architecture has been hand compiled and implemented on a PAL to demonstrate that once a system has been designed, debugged and completed, it can be engineered down to a very small size. Headlights and tail-lights of the toy cars have been replaced by near-infrared proximity sensors.

---

and the can recognizer does not worry about avoiding obstacles. The intelligence lies in the connections between very simple modules. Figure 2 shows Herbert with a network of microprocessors around his torso, the two degree-of-freedom arm extending to the left and a laser light stripe vision system on his head.

### 3.2 Tom and Jerry

Before Herbert was completed, we again digressed and took on another project. This was Tom and Jerry, the automating of two toys cars shown in figure 3. Here, the radio controllers were replaced with a subsumption architecture implemented in combinational logic on one chip, a PAL. The other hardware on the board consists of analog circuitry for the infrared sensors that replace the headlights, and MOSFET switches for controlling the motors. Tom and Jerry again implement three levels of behavior, *avoid*, *wander* and *follow*. [Connell 1987] explains the control system used here in more detail.

The idea demonstrated in the Tom and Jerry project is that the better understood what is being done, the easier it is to engineer it down to a small size. Thus the Lisp machine required as an intelligence engine for Allen has been reduced to a PAL for Tom and Jerry. Another point to be made here is that both the Herbert and Tom and Jerry projects exemplify the flexibility of the subsumption architecture and how, in fact, it is hardware retargetable. This idea is discussed more thoroughly in [Brooks 1987].

#### 4. Future Goals

Over the past two years, we have made a number of observations about most mobile robots and have noted some problems. First, the state of the art in terms of the level of intelligence attainable is not very high. Second, a considerable portion of the sheer bulk on most mobile robots has nothing at all to do with the intelligence system. There seems to be a certain runaway characteristic involved with the sizes of motors and batteries. The tendency in building the stereotypical mobile robot is to start with a chassis and then add a pair of motors and a set of batteries to power the motors. Batteries are heavy, so larger motors are substituted. Large motors draw more power, so bigger batteries are required and so on and on it goes. Meanwhile, the intelligence and sensory systems reside on a few square inches of silicon. Not only does this runaway characteristic of motors and batteries affect the size of the robot, but also the cost. Costs for propulsion systems have not scaled as they have for microelectronics.

Aside from cost and size, there are other factors which have kept mobile robots from becoming as widespread as we might have imagined. Foremost among these is the issue of complexity. A mobile robot is a complex device. Many components such as sensors, actuators, computers and power supplies, usually acquired from separate vendors, must all be interfaced to work together as one system. The more complex such a system is, the harder it is to maintain it. We term these issues the *connector problem*. Much of the sheer bulk on a mobile robot that is not associated with motors and batteries has to do with the connectors, wiring harnesses, cabling and the like. Again, relatively small amounts of space on a typical system are dedicated to sensing or intelligence.

In our line of mobile robots, we have tried to build each succeeding robot smaller and lighter in an attempt to turn around the explosive phenomenon of motors and batteries. An interesting point to ponder is how far down the scaling can continue. It has come to our attention recently that motors are actually being designed on chip, with rotating members on the order of few hundred microns in diameter [Bart, Lober, Howe, Lang and Schlecht 1987], [Trimmer and Gabriel 1987]. If laying down actuators on chip with lithographic integrated circuit techniques is feasible, it may also be possible to integrate the other components of a robotic system: sensors, control and power supplies.

##### 4.1 Gnat Robots

The possibility of creating single chip robots dramatically changes our thinking about mobile robots. Not only would entirely new fields of application open up, but the mass production capability available from integrated circuit fabrication technology holds the promise that someday such robots could be made very cheaply.

Is it feasible? Certainly control logic and microsensors are amenable to such scaling; if power consumption is kept to a minimum, highly efficient solar cells might be usable. But could any type of micromotor develop enough power to locomote a small robot? [Flynn 1987] has taken an in depth look at these questions and has shown that the available power from such micromotors is roughly equivalent to that needed to drive a small vehicle. By combining



such motors with microsensors, a compiled-down version of our subsumption architecture, and miniature power supplies, we can imagine gnat-sized microrobotic systems on a chip at a fraction of the cost of traditional robots, with none of the problems associated with maintenance and spare parts.

Having such robots would change the way we think about robotics and open up new areas of application. Possibilities span from autonomous fishing lures to eavesdroppers to inspectors for spacecraft. Whatever the application, this new model of a mobile robot changes the way we think about using robotic technology. With low cost and small size, we can begin to envision *massive parallelism*; using millions of small simple robots to perform tasks that might otherwise be done with one large, expensive, complex robot. In addition, we also begin to view our robots as *dispensable*. They can be cheap enough that they are thrown away when they have finished their mission or are broken, and we do not have to worry about retrieving them from hazardous or hard to reach places.

To actually build a gnat robot, a research plan is needed that takes the project step by step, from technology that can be produced now to the end goal of what should be developed. Having a clear picture of that goal clarifies what problems should be addressed and in what order.

There are two major guiding considerations here. To create a gnat robot we need to minimize power consumption and save space on chip.

We have a few strategies for design that will enable us to implement these guidelines. First, a gnat robot should rely entirely on passive sensing in order to keep down power consumption. Active sensors, such as those used on our present robots, emit energy to the environment. These types of sensors therefore require significant amounts of power. Although passive sensors usually do not give the information desired directly, we hope to trade off computation for power consumption. Development of better algorithms can keep computational depth shallow and make this feasible (see [Brooks, Flynn and Marill 1987] for an example). A second strategy, one which aids in conserving real estate, is to be clever with sensor design so that one sensor can sense more than one environmental phenomenon. Third, using actuators also as sensors reduces chip size even further. Fourth, we will have to use the most efficient power transmission system possible so that power savings from other strategies are not wiped out.

A plan for the design of a first prototype is detailed in [Flynn 1987] and shows that micromotors can provide power equivalent to that of a rubber band for propelling a small 80mg airplane [Drela 1987]. A research plan is proposed to target the first prototype of a gnat robot to be a single piece of silicon, incorporating sensing, guidance and control systems, which is attached to a macroscopic aerofoil. This example of a small airplane is proposed because it can be the simplest possible system; few actuators are needed and we do not need to worry about micro-aerodynamics and the associated problems of flight in the regime of low Reynold's numbers. This is not giving up on self-propelled gnat robots but is merely good engineering use of abstraction in order to isolate hard problems.

In order to create such a single-chip microrobot for this application it is necessary to pinpoint what problems need to be solved. We can then develop a realistic research plan which best focuses our efforts and expertise to address these issues step by step. Foremost

among the difficult problems involved is the development of micromotors, micro power transmission systems and some means of attaching micromotors to macro linkages. We need more knowledge about the aerodynamics of low-speed, small scale flight. Passive sensors, intelligent control systems integrated down to chip level and small lightweight power sources are other major hurdles. We feel that our most immediate contribution can be in the areas of passive sensing and control systems and it is in these areas where we are focusing our upcoming projects.

All of this is proposed as a direction for future robotics research because we feel that by pursuing this direction, we will eventually be able to mass produce robots and make them cheaply. With gnat robots there will no need for expensive maintenance and spare parts. This will enable them to be widely available and useful even if at first they only have limited intelligence.

#### *4.2 Lunar Rover Robots*

A researcher at the Jet Propulsion Laboratory [Wilcox 1987] has recently thrown out a challenge to the universities. "Build a 10kg lunar rover that can fit into a Space Shuttle Get Away Special, along with its ejection mechanism that can get it to the moon." This would be privately funded, not an official NASA program or anything of that scale. The comment was not made to imply any JPL standing, but to provoke thought and to focus attention on these sorts of issues. The intended degree of effort was that a professor and a few graduate students should be able to produce such a robot in a year.

We feel that gnat robot technology is a perfect fit for this challenge.

Micromotors and microsensors could be used to build even a 10g robot, that could land on the moon, orient itself, extend solar panels, and take pictures which could be transmitted back to an orbiting spacecraft. It could then store some of the energy it collects from its solar panels into a spring, and every so often release the spring and hop to another place to take more pictures. The point is, that if we can make small, lightweight robots, they will be extremely useful for space applications.

The advantage of using gnat robot technology is that hundreds could fit into one Get Away Special, which costs only a few thousand dollars. This would leave room for redundancy and the mission could be declared a success if only one or two ever reached the moon. If such were the case, then less effort would have to be assigned to reliability testing, lowering costs tremendously. A key point is that not only does the small size of a gnat robot open up new options on how to solve a problem, but the potential for *low cost* adds even more value. If the robot is not expensive, there is no worry if it is lost.

There are some interesting issues involved in this lunar rover application. Suppose it were possible to build such a lunar gnat robot. Achieving the moon from the Space Shuttle raises some problems.

Rocket boosters cannot be used to achieve the 7000mph velocity needed to escape Shuttle orbit, because they will not be allowed in the cargo bay due to safety regulations. A rail gun would not be feasible either because it would have to be far too long to fit into a Get Away Special.

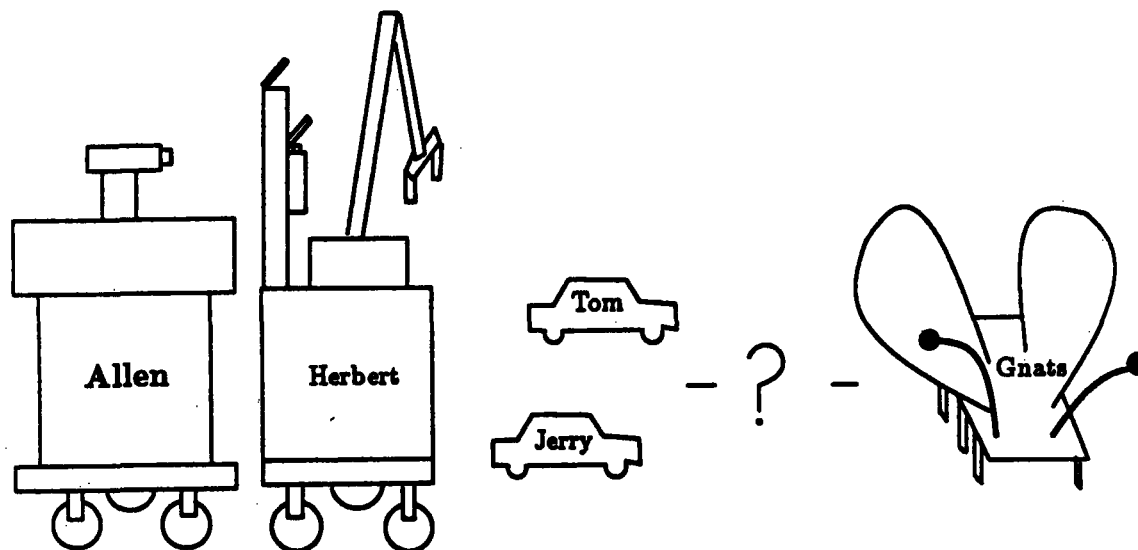


Figure 4. This diagram shows the time sequence of ongoing robot projects: Allen, Herbert, Tom and Jerry. At the right, is the plan that someday we would like to build one-chip, gnat-sized robots.

A colleague, Professor David Akin of the MIT Aeronautics and Astronautics Department [Akin 1987], has suggested a possible solution using solar sails. In the past, all plans for using solar sails were not feasible because the solar sails had to be on the order of a mile square. This was due to the fact that estimated payloads were on the order of 250kg. However, with a 10g robot as payload, one meter square sails would suffice. Sailing to the moon would be accomplished by controlling the sails with micromotors. The robot would need a sensor to sense the moon's outline. One advantage of this whole approach is that the robot could land at a much lower velocity than in the case of being ejected by rocket boosters or a rail gun, in which case it would fall vertically to the moon's surface. With solar sails, it could tack so as to land tangentially and bounce along; not becoming implanted as it would if it just fell. It's a rather fascinating idea; all the more so, because it comes equipped with the features of low-cost, massive parallelism, redundancy and disposableness.

So these are our goals. First, we want to create robots that are more intelligent than the present state of the art. This is manifested in our ongoing Herbert project. We would also like to create small scale robots, with the levels of intelligence we can attain now, in the hopes that low-cost, lightweight, disposable robots will be useful to society and raise our

standard of living. This goal is manifested by our aim to build gnat robots. In addition, gnat robot technology seems to be a fitting and promising technology to be pursued for space applications.

## 5. The Research Plan

We have spent a few years building some prototype robots and gaining working knowledge and experience. We have outlined what we see as hard problems and directions that should be pursued. We also have some lofty goals. Figure 4 shows a time chart of the sequence of projects which we have begun and also the goal projects we would someday like to build.

How do we form a research plan that successfully leads a project step by step to meet these goals? We want to show that we are not talking about science fiction, but mean to achieve these goals in the not so infinite future.

In order to identify the steps needed to bring this to fruition, let's work backwards. Starting from the desire to have a working, flying gnat robot, what would be the next easiest project, one step back that would be doable without facing all the problems involved in gnat robots?

### 5.1 *Micro Robotic Airplane*

[Flynn 1987] proposes building a miniature airplane. Basically, the suggestion is to put together an entire robotic system on a chip and then mount that chip on an airfoil which requires only one actuator for power and one actuator for control. What pieces would have to be developed to realize this? A passive sensing system would have to be built so that the airplane would not crash. A control system would have to be laid out on an integrated circuit to turn the sensed data into appropriate actuator commands. Micromotors would have to be built. Complex, highly efficient solar cells would have to be produced. Some type of transmission system would have to be conceived to couple the micromotors to a macro propeller or rudder.

The hardest problems here are probably the micromotors and solar cells. However, we feel we have a good angle of attack on the sensing and control issues. Our aim therefore is to concentrate on developing the technologies in these areas, and to hope for outside progress in micromotors and solar cells. Advances in micromotors and solar cells will be brought about by those with more expertise. Our strategy then is to take one more step back and aim towards a project that isolates the problems where we feel we can make a contribution.

### 5.2 *Autonomous Rubber Band Powered Airplane*

Consequently, a more immediate stepping stone would be to take the same lightweight airplane and leave the rubber band on to power it. This would give roughly six minutes of flight. Automating that airplane for that time interval would still be quite remarkable.

We will work on building a miniature vision system that provides input to a compiled down subsumption architecture. This sensing and control system will give the airplane the

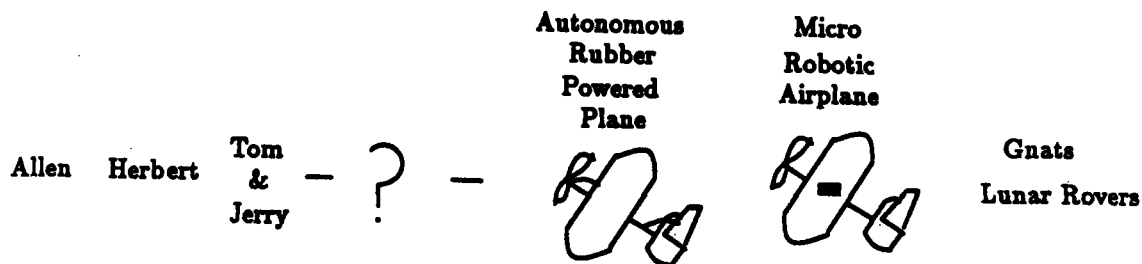


Figure 5. Shown here are the projects to aim for that are just simpler and easier to attain than gnat robots, but develop parts of the technology needed. A microrobotic system built on a chip and then attached to a macroscopic aerofoil is one step back from a true gnat robot. However, doing that is still quite involved, so automating a rubber band powered airplane would be a project that would require implementation and integration of a small sensing system and a compiled subsumption architecture powered by solar cells, without the need for micromotors.

---

ability to fly around, avoiding obstacles, turning left or right as appropriate. The hand wound rubber band can remain, in order to provide power. Control of the rudder could be performed by a small nitinol actuator [Trimmer and Gabriel 1986] or possibly even a scaled down piezoelectric actuator. Figure 5 shows how these projects fit into our overall research plan.

All of this leads to the idea that what our mobile robot group should be focusing on today is creating the sensors and control systems necessary for gnat robots. Of course, we are also still interested on the parallel track of making our present robots exhibit more intelligent behavior. This leads us to our present plans for our next few projects.

### 5.3 Seymour

In order to pursue some of the directions we have been discussing, we have decided to start a new project. This new robot, Seymour, will specifically be targeted at developing passive sensors. We are going to constrain ourselves, artificially, to these low power devices in order to see how much sensory information we can extract and how intelligent a robot we can build, based on passive sensors. In addition, although Seymour will be a macroscopic robot, we will try, as best we can, to keep power consumption down and to make him as small as possible. At the moment, our plans are that he will be 12 inches in diameter and 20 inches high.

To go to all the trouble of beginning a whole new project, we should plan to get as

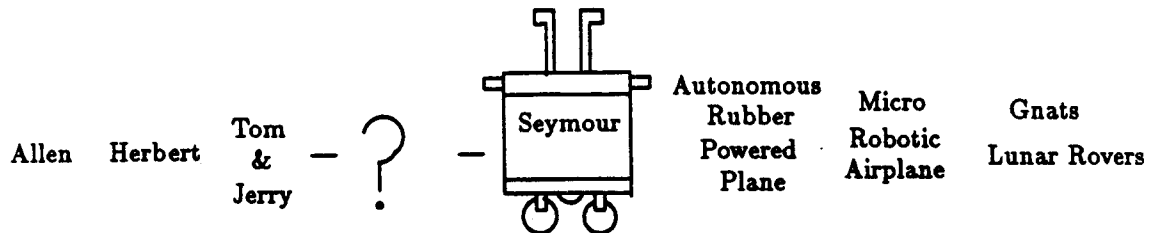


Figure 6. Seymour is the robot on the drawing board at the present time. Constrained to move around only through passive sensing in order to push the technology required for gnat robots, Seymour will also be able to exhibit increasingly complex behavioral patterns since he will have many mechanisms for dealing with the outside world. Namely, he will be a Business Bot, wandering around selling candy, and hopefully clever enough that he will be able to improve himself and become more successful at his trade as he gains experience.

much out of it as possible. Therefore Seymour will be a workhorse not only for pushing towards gnat robots but also for pushing the levels of intelligence which we can incorporate into our robots. Building an entirely new robot requires a lot of work. A new base has to be built, power supplies have to be installed, computers have to be interfaced, and so on. Since Seymour is very small, there is little room to slap things together or throw things on later. Very careful system design is required from the start. Although we have quite a number of engineering reasons for building Seymour, such as pushing the technology of passive sensors, low power consumption and small size as far as we can, in order to make progress towards gnat robots, we might as well also work on the parallel track of making our robots more intelligent. We will try to do this however, within the engineering constraints which we have set for ourselves. So Seymour will be more than just another robot to avoid obstacles, but more will be said about that later. Obstacle avoidance using passive sensors is a hard problem in itself. Figure 6 depicts how the Seymour project fits into the research plan. Although Seymour is a macro-robot, it is a first and very real step towards building small scale robots.

### 5.3.1 Stereo and Motion Vision

The constraint of using passive sensors is a tough one. One can acquire direct information and acquire it much more easily by using active sensors such as light strippers, infrared proximity sensors or sonar sensors. We are hoping however, that by coming up with clever

algorithms for passive sensors such as CCD cameras, we can trade off power consumption for computation, as we can do computation with low-power CMOS microprocessors.

One of our first goals on this project is to get Seymour to perform navigation, that is, avoid obstacles in real time using cameras and onboard vision processing. A major problem in trying to use vision is that most vision algorithms don't work on real data or are not particularly useful for robot navigation. This is in spite of the fact that biological visual systems provide an existence proof of the feasibility, robustness and richness that a vision system should be able to provide.

Machine vision systems require extensive computation and more research needs to be done to develop algorithms which can cope with all the noise and complexity of visual perception. We have recently been working on this problem and have developed some new algorithms which provide useful, reliable information to an autonomous robot for the purposes of obstacle avoidance [Brooks, Flynn and Marill 1987].

The idea is to use two early vision processes, stereo and forward motion analysis, to determine how far away obstacles are. Stereo and motion processes are very sensitive to misalignment of the cameras however, so the key component of this new algorithm is that it is self-calibrating. The idea of self-calibration is to have the stereo and motion routines cross-calibrate themselves off the real world. This is what really makes the algorithms work.

One benefit of self-calibration is that no test patterns are needed for the robot to stare at upon being powered up, but rather the robot calibrates itself by stumbling and blundering about the real world and gradually becomes able to see better as time goes on. It runs its calibration routines in background continuously and should someone bash one of its cameras and knock it out of alignment, it soon figures out the new parameters. In this way, the robot is able to predict reliably time to collision with various obstacles and so determine the correct motor commands in time to avoid them. All of this is done passively and therefore consumes relatively little power.

The stereo and motion algorithms keep computational requirements within the bounds of onboard processing by using single scan line cameras with cylindrical lenses. These lenses average the image optically and highlight vertical edges. The algorithm finds edges and then tracks them as the robot moves. If a camera is misaligned slightly from the direction of motion, the center of expansion will not be in the center of the image. Its location can be determined however, by finding the point from which the edge trails diverge over a sequence of images. If the speed of the robot is constant, then time to collision to various edges can be calculated.

If this motion algorithm is used alone, then no information can be obtained when the robot is stationary or turning. Also, no information can be obtained in the field of view directly ahead. For these cases, we use stereo. The usual problem with stereo is that if both cameras are not perfectly parallel, serious errors in calculated range are returned unless the cameras are calibrated. In most cases, a calibration routine is run when the robot is powered up by having the cameras stare at a known test pattern. In our algorithm, we forego test patterns and initial calibration by using the motion edges to provide roughly 30 good matches for the stereo, where we have depth estimates in both the left and right

cameras. These 30 matches are then used to run the least squares fit in the stereo calibration routine.

We have tested these algorithms on Allen using array cameras connected to a Lisp machine. The images were averaged down to one line and all the stereo and motion code ran on single scanlines. The algorithms seemed to work in that they reliably predicted time to collision to various obstacles as we manually drove the robot. However, the Lisp machine was not fast enough to grab images and do the calculation in real-time. For this reason, we are building Seymour, which will not only have real line scan cameras with cylindrical lenses but also an onboard vision processor which we are designing in-house. (Seymour is going to "see more" and therefore will "seem more" intelligent.) The vision processor consists of four boards; one 68000 board, one board with an Analog Devices digital signal processor, one memory board and one camera interface board. Each board is roughly 4" x 8". Whereas the Lisp machine could grab roughly 15 frames a second, this system should be able to grab over 1000. We plan to run at a 10Hz rate and use the remaining time for processing and for other vision algorithms running on the same inputs. For instance, we may have side looking single scan line cameras to identify doorways and corridors, forward looking small arrays ( $32 \times 32$ ) to look for overhead lights along corridors, and steerable cameras for object tracking.

The vision system we have so far demonstrated is the appropriate sensor for mobile robot obstacle avoidance. Since we are not trying to recognize the obstacles we are avoiding, we can use linear array cameras and that makes it possible to build a processor which fits onboard. Also, the algorithms deliver the right type of information, because instead of returning distances to objects which would have to be mapped into some absolute coordinate system, they return time to collision. The control system can then simply issue velocity commands to the robot to make it move out of the way. This type of vision system also is a nice fit with our desire to build autonomous robots, in that the self-calibration feature allows the robot to become autonomous upon power up.

### *5.3.2 A Cohesive Theme*

Taking on a new robot project involves an extended effort. Even though real time visual navigation would be quite a feat, we would also like Seymour to have a goal in life that requires more sophisticated behaviors than obstacle avoidance. The project has basically taken its shape due to the types of problems we want to address: passive sensing, development of better visual algorithms, low power consumption and small size. However, if we could stay within these constraints and get even more out of our project, we would save a lot of effort. Just as Herbert has a goal in life of collecting soda cans, which allows us to address the issues involved in building more complex behaviors, Seymour too needs a goal in life. We searched to find a testbed in which, if we just added a few more simple mechanisms for Seymour to deal with the real world, we could have a fairly rich scenario for researching somewhat higher levels of intelligent behaviors.

The scenario we converged on is that Seymour will be a Business Bot. Seymour's goal in life will be to make money. Basically, Seymour will accept payment for services rendered.



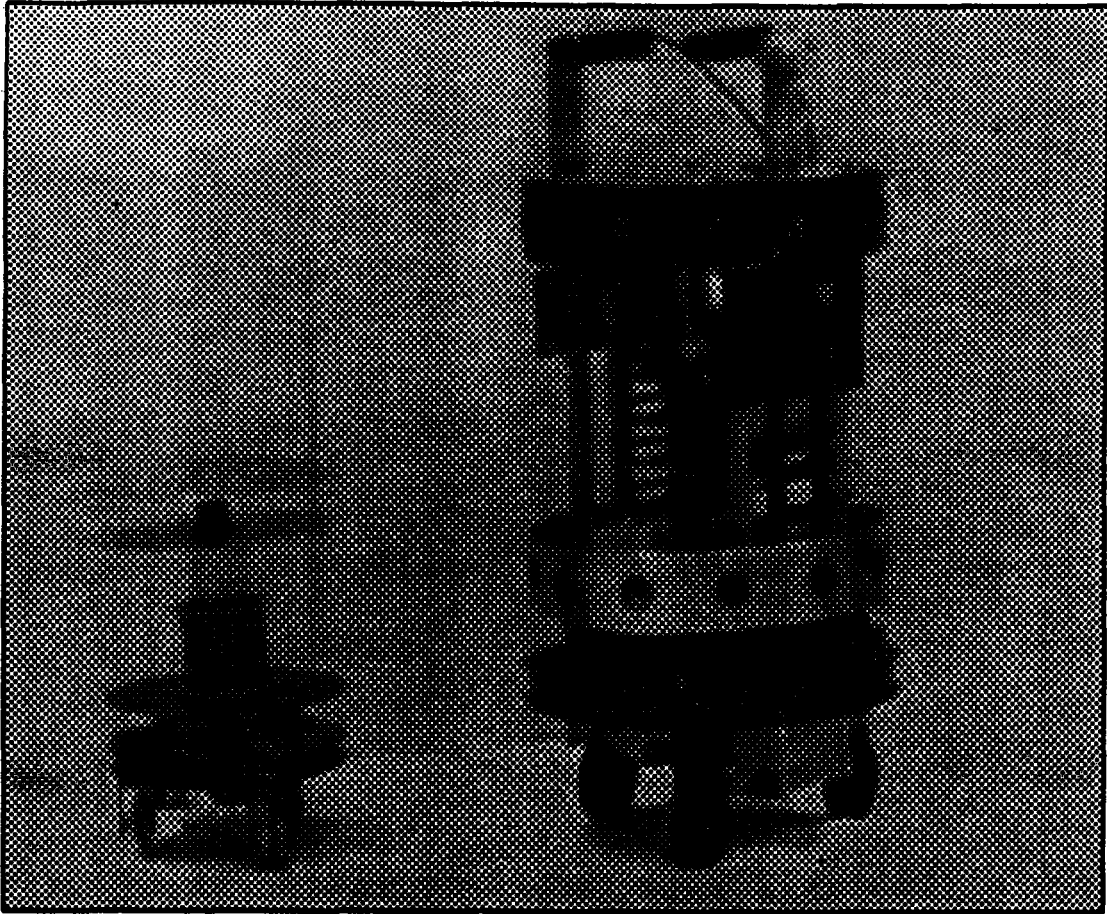


Figure 7. Seymour will be much faster and smaller than either Allen or Herbert. Shown above in proportion to Allen, it's clear that the engineering goals for Seymour are ambitious, as Seymour will have far more sensors and computational power than Allen ever had. At the present time, the base is all that is complete. The upper structure is a mock-up to visualize this final diameter and height.

Pay him money and he will give you a song or dance, tell you a joke or sell you some candy. Then with the money he makes, he pays people for things he wants: the door opened, the elevator button pushed, his candies restocked, maybe even a tune up.

This scenario was picked because we wanted to develop an intelligent robotic system without spending too much time or energy on complex perception or action systems. We claim that robotics is the intelligent connection of perception to action, but years of work have been spent on the very hard problems of vision and manipulation and few people have gotten around to working on any intelligent connection. Seymour will help bridge that gap, by becoming known as a *clever* robot while not being especially *perceptually sophisticated*

or *manually dextrous*.

### 5.3.3 *The Consumption Architecture*

All of the behaviors for the candy selling robot will be controlled in the same distributed manner as the previous robots. We nickname Seymour's control system *the consumption architecture*. The vision system will provide sensory input for the lowest levels of the control system, just as sonars did on Allen and infrareds did on Herbert. Higher levels for selling candy will again be triggered by various sensory events. A typical scenario is that Seymour will randomly wander around, avoiding obstacles if necessary. If a person is detected, an approach behavior is triggered. When he gets close to the person, he starts talking: "Wanna buy some candy, good price, special deal just for you, twenty five cents...". If he senses money input to his coin recognizer, he dispenses the candy, and so on.

Traditional robot projects often deal with the problems of recognition and manipulation. Seymour thinks about recognition and manipulation differently. Instead of recognizing shapes or geometric objects he will be more interested in recognizing situations and aggregate behaviors of people. Instead of manipulating objects and geometric shapes, he will manipulate people—with a whole bag of tricks for making money for selling them things and getting them to do things for him.

### 5.3.4. *Building Seymour*

Seymour is still on the drawing board, but most of his pieces are already in hand and he is coming together as this is being written. Figure 7 shows the present state of progress. Physically, Seymour will be built into a base similar to our earlier models, only this base will be smaller and much faster. Bumper sensors will protrude about 1/2", tripping microswitches if actuated. The brains will consist of a 68000 based vision processor with a digital signal processing chip, and a 68000 based consumption architecture. Linear cameras with cylindrical lenses and passive infrared motion sensors will be the primary sensors, to be used for avoiding obstacles and following people respectively. Although these sensors and the computer boards are being designed in house, many components have been purchased off the shelf. A simple speech synthesis board is used for sales pitches. A flux gate compass will assist in navigation and a real time clock will be used to keep track of transactions. Business is transacted with simple mechanisms for dispensing candy and recognizing coins. A coin dispenser is also mounted onboard so Seymour can purchase services. A 640 x 200 element liquid crystal display connected to a small keyboard will rest on top to provide interaction with customers and to assist in debugging. This will be especially useful for displaying the edges found by the cameras.

### 5.4 *Silicon Compiler*

Concurrently, we have started a project to build a silicon compiler and custom VLSI implementation of the subsumption architecture. The idea is to write a program which takes

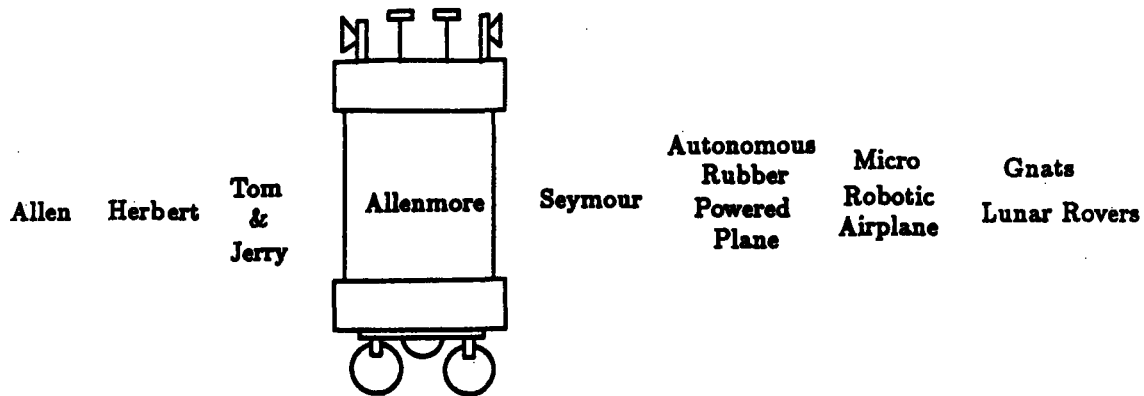


Figure 8. Allenmore is a prototyping project for subsystems that will go into Seymour. Specifically, we will build a prototype line scan camera system and passive infrared sensors with ranging capability and test them out on Allen. Because Allen's subsumption architecture is so flexible, due to the fact that it is implemented on an offboard Lisp machine, we can quickly test out these new sensor designs and see how well they work within the planned control system, without waiting to have all of Seymour finished.

the control system specifications, which are written in Lisp and are basically finite state machines, and macro expands the code into logical building blocks. An optimizer then takes that level of description and compacts it to the minimum number of gates required. This gate level description is then used as a netlist and VLSI implementation is taken from there. The plan is to eventually compile the entire control system onto one chip.

### 5.5 Allenmore

A good deal of engineering is required to develop the sensors we need into package sizes appropriate to fit on Seymour. In addition, we need to prototype and test our ideas for the sensor and control algorithms. Fortunately, we can use Allen as a testbed robot to do some quick implementations of these ideas. Once the functionality is established, the sensor can be carefully designed and packaged to fit into Seymour. This prototyping project for Seymour we term Allenmore. Figure 8 depicts how the Allenmore project fits into the timeline of projects.

At the moment we have three development projects ongoing with Allenmore. First, we are building a system which will allow us to interface the linear cameras into a Lisp machine so that we can test our algorithms with the real cameras that we plan to use on Seymour. Coincidentally, the cameras are interfaced to a DSP board which resides in a

Lisp machine. That board has a port on it through which we can grab linear images very quickly, transfer them to Lisp arrays and run our Lisp algorithms which we had tested earlier with 2D-array cameras. Although the images from those cameras were fed into the Lisp machine through the frame grabber over a video cable, we could not use that system to grab linear images. After experimenting with the linear cameras, we can if we desire, rewrite our Lisp algorithms to run on the DSP card. Although the DSP chip in the Lisp machine is from Texas Instruments and the one on Seymour is from Analog Devices and the code will not be directly transferable, having such an interface to a Lisp machine lets us easily do initial development for new algorithms that we might later want to retarget to Seymour. By running vision algorithms on the DSP card, the rest of the subsumption architecture can still run on the Lisp machine and control Allen's drive motors. This way we can see how well our vision algorithms work in connection with the control system – something we have not yet been able to test.

A second project being undertaken on Allenmore is the prototyping of the passive infrared motion sensors which Seymour will use for detecting people and possible markets. These pyroelectric sensors respond to a change in temperature across their field of view. We are experimenting with different configurations of these sensors in order to see what works and how we can best implement a person-following behavior. It would be nice to be able to get some range out of these sensors, but at the present time they provide only a binary output of whether or not they detect any change in temperature. We are looking into the possibility of using geometry and convergent beams to do crude ranging. Another idea is to spin two sensors, threshold the analog output for peaks in temperature change, and then do stereo triangulation. The output of these sensors can be sent over Allen's cable to the Lisp machine and we can extend the subsumption architecture there to include the control for a people-following behavior. Having available one robot to use as a testbed for developing subsystems is a tremendous help in getting the new robot up and running.

In addition to doing prototyping work for Seymour, we have a third project in the works to extend Allen's control system for its own sake. Basically, we started by building three levels of behavior for Allen but then digressed and took on several other projects, and the question is often raised as to why we have not added more layers. We are presently in the midst of experimenting with two more layers using the sonar sensors already mounted on his top. In addition, we have recently added a second ring near the base, which should provide better coverage. The new layers of control will use both of these rings to implement some corridor finding and wait-at-the-end-of-the-hallway behaviors.

## 6. Conclusion

This paper has tried to explain the goals of our project and our methods for actually achieving those goals. We have shown, step by step, the projects we have planned and how we plan to develop each succeeding stage of technology. By discussing all of our past, present and future projects in the context of what we are trying to achieve, we hoped to more clearly explain why we are working on our present projects. We hope that someday,

this line of research will lead to mobile robots which are both more intelligent and more useful.

## References

- [Akin 1987] David Akin, MIT Aeronautics and Astronautics Department. Personal communication.
- [Bart, Lober, Howe, Lang and Schlecht 1987] "Microfabricated Electric Actuators", S.F. Bart, T.A. Lober, R.T. Howe, J.H. Lang and M.F. Schlecht. To appear in *Sensors and Actuators*.
- [Brooks 1986] "A Robust Layered Control System for a Mobile Robot", Rodney A. Brooks, *IEEE Journal of Robotics and Automation*, RA-2, April, 14-23.
- [Brooks 1986b] "Achieving Artificial Intelligence Through Building Robots", Rodney A. Brooks, *MIT AI Lab Memo AIM-899*, May.
- [Brooks 1987] "A Hardware Retargetable Distributed Layered Architecture for Mobile Robot Control", Rodney A. Brooks, *Proceedings IEEE Robotics and Automation, Raleigh NC*, April, 106-110.
- [Brooks 1987b] "Intelligence Without Representation", Rodney A. Brooks, *Proceedings of Foundations of AI Workshop, MIT*, June.
- [Brooks and Connell 1986] "Asynchronous Distributed Control System for A Mobile Robot", Rodney A. Brooks and Jonathan H. Connell, *Proceedings SPIE, Cambridge, MA*, October, Vol 727, 77-84.
- [Brooks, Connell and Flynn 1986] "A Mobile Robot with Onboard Parallel Processor and Large Workspace Arm", *Proceedings AAAI Conference, Philadelphia, PA*, August, 1096-1100.
- [Brooks, Connell and Ning 1987] "A Fully Autonomous Subsumption Architecture Mobile Robot", Rodney A. Brooks, Jonathan H. Connell and Peter Ning, *in preparation*.
- [Brooks, Flynn and Marill 1987] "Self Calibration of Motion and Stereo Vision for Mobile Robot Navigation", Rodney A. Brooks, Anita M. Flynn and Thomas Marill, *MIT AI Lab Memo AIM-984*, August.
- [Connell 1987] "Creature Design with the Subsumption Architecture", Jonathan H. Connell, *Proceedings IJCAI-87, Milan, Italy*, August.
- [Drela 1987] Mark Drela, MIT Aeronautics and Astronautics Department. Personal communication.
- [Flynn 1987] "Gnat Robots (And How They Will Change Robotics)", Anita M. Flynn, *IEEE Micro Robotics and Teleoperators Workshop, Hyannis, MA*, November.

[Lang, Schlecht and Howe 1987] "Electric Micromotors: Electromechanical Characteristics", Jeffrey H. Lang, Martin F. Schlecht and Roger T. Howe, *IEEE Micro Robots and Teleoperators Workshop, Hyannis, MA*, November.

[Trimmer and Gabriel 1986] William Trimmer and Kenneth Gabriel, AT&T Bell Laboratories, Holmdel, NJ. Personal communication.

[Trimmer and Gabriel 1987] "Design Considerations for a Practical Electrostatic Micro Motor", W.S.N. Trimmer and K.J. Gabriel, *Sensors and Actuators*, 11(2):189-206.

[Wilcox 1987] Brian Wilcox, Jet Propulsion Laboratory. Personal communication.