MASSACHUSETTS INSTITUTE OF TECHNOLOGY
A. I. LABORATORY

Artificial Intelligence
Memo No. 251

January 1972

MINI-ROBOT PROPOSAL TO ARPA

Marvin Minsky

This proposal was written at the Artificial Intelligence Laboratory, a Massachusetts Institute of Technology research program supported in part by the Advanced Research Projects Agency of the Department of Defense and monitored by the Office of Naval Research under Contract Number N00014-70-A-0362-0002.

Proposal for Research on Micro-Automation

Marvin Minsky, M.I.T.

0.0 Introduction

During the next decade it will become practical to use more and more sophisticated techniques of automation -- we shall call this "robotics" -- both in established industries and in new areas. The rate at which these techniques become available will depend very much on the way in which research programs are organized to pursue them.

The issues involved are rather large and touch not only on technical matters but also on aspects of national economic policy and attitudes toward world trade positions.

The project herein proposed is concerned with the development of two particular aspects of Robotics, namely:

1. Development of a miniature hand-eye system.

2. Development of remote, ARPA-NETWORK style operation of robotic systems, in which simple jobs are handled locally while more complex computations are done on larger, more

centralized, machines.

The following sections explain why we feel that these areas deserve particular attention at this time. We will discuss some general issues first, and then concentrate on the particular subjects noted above.

1.0 Automation and Mass Production

The great development of mass production, in the early twentieth century, centered around the realization that machine tools could be made so precise -- without prohibitive expense -- that one could build stocks of interchangeable parts for later assembly by unskilled labor.

This led to a huge variety of automatic production machinery specialized to fabricate parts and components. A similar development took place in the manufacturing of materials to meet new standards of uniformity imposed by the needs of the automatic fabrication equipment. The modern assembly line took shape. But the Automation of Assembly itself did not develop so rapidly, and today it is the major consumer of unskilled labor in the factory.

The direction of progress in this general development was always toward precision and inflexibility. The integrity of the production line increasingly came to depend on uniformity of materials, precision of positioning and repeatability of the machine processing operations. The major investment in setting up such assembly lines restricted their use to high-volume operations.

In the late 1950's some new elements began to appear in the automation area. "Numerical-control" production equipment came to maturity. A numerical-control machine is, more or less, a conventional machine tool such as a lathe or a milling machine, with motorized controls that are operated by a preprogrammed recording (usually on a punched tape). Such machines can produce very complex machined parts economically because the extra costs of such machines pay off in reliability and speed over manually operated machines (where a single mistake can ruin hours of work). The skill of the human operator, in making judgments about quality of cut and compensation for tool wear, is replaced by careful planning, control of materials, and scheduled replacements of tools. In sophisticated systems, such parameters as tool wear compensation are programmed in advance.

Along with this came some advances in automatic inspection. Devices have appeared which can verify, afterward, whether the preprogrammed operation acheived its goal. For example, one can buy machines that can be programmed to probe the walls of an engine block to see that the cylinders were properly bored, that the interior surface is intact, and that the dimensions are correct to within required tolerances.

It would be mistake, however, to get the impression of a continuous, rapid, and steady progress in automation. The modern factory is not very different from that of twenty years ago, and there is little backlog of technical ideas for further advances, because research in advanced automation has come to a near standstill in recent years. Except for the projects sponsored by the Advanced Research Projects Agency, and work in Japan, we know of no sophisticated industrial research projects in this field. Certainly, we know of no such projects in any industrial research laboratory!

In particular, fabrication and inspection have not been very well mated, on the whole. There is very little "cybernetic" activity in modern factories; few material-handling processes use advanced kinds of feedback interaction. Because of this, we find little automation of complicated assembly, and no significant work on automated service, maintenance, or repair.

The outsider gets a different impression. "Industrial robots" have appeared on the public scene, in the form of the pre-programmed positioner-grasper machines of the Unimate or Versatran class. But except in a very few cases, these are used in much the same manner as special-purpose positioners, without significant sensory feedback. Because this lack of "intelligent feedback" makes it impossible to use the potential versatility demonstrated in the ARPA robotics experiments, these "robots" have not yet had any really substantial impact on production.

To summarize: we have the impression that despite a superficially healthy appearance, the past two decades have seen a pernicious decline in acquisition of basic new production techniques. What is worse, there is little indication that the community of economic advisors has sensed this technical aspect of the productivity problem.

2.0 New Application Areas

A vigorous research and development program, over the next decade, could lead to substantial progress in such areas as listed below. These are only examples of applications; it is hard to see how any field could remain unaffected.

2.1 General Production

The outstanding applications of Advanced Automation to production, in the near future, will be in assembly, testing, finishing, etc., rather than in the fabrication of large numbers of simple parts (which is already substantially automated) or even complicated parts, like engine blocks (which also are economically produced by special purpose systems). But advanced automation will nonetheless profoundly change the factory because of these possibilities:

Assembly lines

The traditional special-purpose automation devices were based on having just a few operations at each "station". This dictated the linear arrangement of work stations, with the work moving from one point to another, and different parts brought to different stations. The general-purpose assembler would work more centrally, with many operations performed at the same station, and the parts all brought to that point. If the assembly machines can SEE, the parts need not be so carefully positioned and could come in mixed arrangements, on demand, or on prepared trays.

"On Demand" production

The arrangement just described might seem at first sight a step backwards! Would not the centralization of assembly (instead of distribution along a line) cause traffic problems? This potential problem could usually be resolved by mixing parts on a very few supply lines. And the payoff would be large for these indirect reasons: To begin with, the versatility of the system reduces the length of the production "runs" required for economy. Today, every run requires the committment of substantial space and effort to set up the line and construct the special purpose machines at each station. This is too expensive unless the run produces a large stock of items. Accordingly, one can only attempt to balance the anticipated demand against the expenses of carrying large stocks and tying up capital investment. In any case, one has to pay the price of long advance order times. With the advanced system one will be able to switch from one product to another without stockpiling many finished items. One still needs to have components available, to the extent that they are not also subject to "on demand" fabrication.

Preparedness without stockpiling

The attractive possibility, then, is to be able to get rapid production of complicated units quickly when they are needed, without pre-production of all the items one might need quickly under all significant contingencies. This also means that in anticipating a future demand, one need not make, in advance, uncertain specifications! Instead, one may be able to write conditional branches into fabrication and assembly programs, so that one can make the final specifications just before delivery time.

"Individualized" mass production.

Perhaps the most exciting aspect of Advanced Automation, from a social and humanistic standpoint, is the way in which it can reverse the "mechanization" that was a natural consequence of the first industrial revolution! The inevitable result of early mass production was uniformization of products. Clothing, for example, had to come in a standard variety of sizes. Shoes were characterized by two parameters, "width" and "length". While the production economy meant that there could be more shoes, they had to be worse fits! So the quality of life had to give way to the quantity. This price no longer needs to be paid!

The new machines could be much more responsive. Clothes could be cut and made to fit a dimension file (or even picture!) for each person. Plumbing modules could be fabricated for each particular installation, given a file describing the current state of a house-plan. We leave this whole area to the reader's imagination, simply noting that while one can not expect much change in the costs of materials, over the next few years, one can be certain that the cost of computation will continue to drop.

2.2 Medicine and Health

Anyone who has been hospitalized will remember enough complaints to make enumeration of examples here unnecessary. We recognize that there has been rapid progress in semi-intelligent data processing for medically relevant analytic and diagnostic procedures. (Even more has been done to automate, for better or for worse, the fiscal operations of the hospital.) But, except for some procedures in the hospital laboratory, there has been no improvement in the quality of personal patient treatment and care.

There are too many factors contributing to the inadequate, degrading, expensive, and downright dangerous aspects of life in the hospital for us to review here. But it is clear that many of the most serious problems stem from the economic impossibility of providing continuous, sensitive, and unobtrusive attention to the patient's state -- medical parameters, physical position, mobility, and comfort, and his needs and wants -- to say nothing of adapting his treatment schedule to his actual (rather than his anticipated) condition.

We can best show how advanced automation is relevant to this area by specific examples. An immediate application of physical robotic techniques is needed in Radiation Therapy, an area of medical treatment that is currently in such short supply, and so poorly administered, as to constitute a public scandal and a continuity of private tragedy.

In attempting to destroy a region of tissue, by appropriate forms of radiation, one wants to deliver carefully calculated doses of radiation to the site. Because other structures between the area in question and the beam source must get smaller doses, elaborate computer programs for treatment planning have been written and debugged that integrate many small doses from carefully positioned beams in many directions.

The problem is not simply geometric; one may want to concentrate the dose on tissues close to the spinal cord, or to the aorta, which may be exceptionally sensitive to smaller amounts of stray exposure. The calculated treatment plan may appear to be adequate, mathematically, but delivering the dose is another matter. Patients move. Beams are shaped by a combination of aperture shaping at the beam source and stacks of lead objects near the patient. Setting these up takes time, and they may slip, etc. Ideally, one wants to move the source through a smooth trajectory, with corresponding motion of the shielding blocks, but this is impractical today, so we approximate by a few fixed positions. In any case, it is extremely hard to be sure that the patient position stays within the reqired few millimeters. In a very few centers in the world, expensive machines and a substantial, highly trained staff, can give an accurate treatment in thirty minutes or so. (A course of treatment is a great many such sessions.) Even so the dosage is uncertain and accidents of positioning and exposure are much more common than anyone will dare admit. But 99% of the people needing such treatment do not get to these centers, and do not even get the benefit of the mathematical precalculation of the proposed dosage.

The solution is evident: the beam position, aperture, and shielding set-up should be automated, as should the manipulation of the patient's gross position. Visual monitoring of reflecting markers on the patient's surface should be fed back to the machine's physical adjustments, and to its beam modulation devices. Fluoroscopic feedback from the patient's internal structures could also be processed by the vision computer, using natural structural shadows or implanted radio-opaque markers. Files of the progress of the treatment plan could be maintenance in the system, and the exposure built up accurately. The effective duty cycle of the machinery and staff would thereby be raised to a reasonable level.

Other medical applications

The applications of miniature robotic hands in surgery are simply fantastic, even if used without any computer at all, simply with remote master-slave control. One could do extensive internal examination through a very small laparotomy without much trauma, and many kinds of repair. One could do mechanical engineering in the middle ear, the heart, the eye, dentistry, etc., etc. The domain of microsurgery could easily be greatly extended. We consider it best to leave such visions to the imagination, because spelling them out could only arouse skepticism. People who have to read this proposal through bifocals will agree that our great technology ought to be able to offer them more help. Imagine an adjustable eyeglass lens activated by a sensor that is operated by the eyebrows, or better, by a tiny sensor buried in the muscle within the eye that is unsuccessfully trying to focus the ocular lens.

2.3 Mining and Undersea Resources

No one should have to be exposed to dangerous mining conditions. The elaboration of versatile inexpensive master-slave devices, again even without advanced automatic control systems, should totally change this situation. Deeper, narrower, more extensive subterranean exploration ought to open all sorts of new resources to economic development. The same obviously holds for undersea activity, and undersea mining. There is no reason why mining should not become an elegant white-collar activity, carried on from perfectly comfortable control sites, even one's home!

Indeed, the fact that all sorts of skilled manual labor could be done from arbitrary locations should have a deep, and generally positive, effect on many aspects of our crowded city life! It is fun to imagine the full consequences of having versatile, economical, remote master-slave devices.

2.4 Housing

Housing is also subject to improvements both in fabriction and assembly. The current concept of "prefabrication" is a very weak form of the idea, for houses ought to be adapted for the needs of their inhabitants. The design phase, with substantial programmed help, should produce programs both for fabrication of the parts and for assembly at the site, taking the special features of the site into account. The machines used today for site preparation and construction are already, in fact, more advanced as manipulators, than the equipment used in factories; and we can expect them to move further in that direction. It is difficult to predict very much of what will happen in this area, of course, because it is so subject to radical changes in materials and methods.

2.5 Transportation

Automation of the transportation of people has been extremely slow to develop, presumably because of the extreme constraints on safety. Imaginative proposals of all sorts, from simple traffic control devices to advanced moving belts and capsules, have always been rejected because there was no technology for reliably sensing where the people are and what they are doing. The possibility of eyes that can really see, and hands that could swiftly but gently correct a dangerous situation, could make a great difference in this area.

2.6 Space

Advanced automation could have a dramatic effect on the character of space exploration and application, by drastically changing the balance of costs and reliability. There are three basic factors here.

1. The use of a central manipulator-sensor system to operate all onboard systems can greatly increase the effective payload of a mission, by reducing the gross weight of the experiments and equipment.

2. The use of the manipulator system makes possible the remote maintenance and repair of spacecraft systems, removing redundancy, and using the same structural parts for different applications.

3. Power, size, weight, structural strength, endurance, process speeds, reliability, maintenance and hazard control all become less rigidly constrained without the astronaut, yet mission objectives can still be sophisticated and adaptable.

The basic idea is simple. The reliability of the mission should be invested primarily in the central computer and the hand-eye system. Then it becomes possible to diagnose faults and correct them by changing modules or making minor physical repairs with EVA and zero G manipulation being mainly programming problems. Even faults in the motors moving the eye can be bypassed by using the hand to make such adjustments.

The versatility of the general-purpose hand-eye system can be exploited in the construction of very large, delicate, antenna systems in space, avoiding the requirements of complex selferecting structures and, possibly, reusing materials no longer needed in earlier phases of the mission. Increased antenna size reduces the communication power requirments and hence the overall weight of the payload. The same is true for activity at a remote landing site, such as laying out a phased array antenna on the ground.

Another important point concerns internal measurements. Ordinarily, to measure the temperature of, say, a fuel tank, one would implant a sensor and run a wire to a central telemetering station. If one imagines a camera eye in the spacecraft cabin, one can simply weld a bimetal appendage to the surface in question, and view its deflection visually with the eye. At virtually no cost in equipment or reliability, one could mount thousands of simple sensors to any number of surfaces within the spacecraft, measuring temperature, pressure, strain, current, etc., etc., without any central connections, reading them all with the eye.

Similarly, the scientific and navigation instruments could be loosely coupled to the system. Even celestial observations could be made by the eye, through convenient lenses and reticles, without complicated special instrumentation. Dials on communication equipment could be read, and simple adjustments made by hand, trimming antennas, deflecting vernier thrust units, etc, etc.

It is our impression that the philosophy of spacecraft mission design has evolved continuously in the direction of "fail-soft" reliability engineering, making subsystems more and more autonomous so that some will function when others fail. This, we feel, is a typical evolutionary trap that one cannot break out of by small steps! We conjecture that the time has come to consider the possibility that the chances of mission success are greater if we do the opposite -- put all our eggs in one basket! -- namely, in the central-computer hand-eye system -- and use its versatility to provide the fail-soft quality of the mission in a centralized way. If the computer fails, the mission is lost, but we have come to accept that this is true for the booster, and should see it in the same light.

Typical of the resulting payoff: in space, a few simple beams and adjustable clamps, a modest variety of optical and electrical components, etc., make it possible to construct an inexhaustible variety of optical instruments, spectrographs, analytic devices, image transformers and filters, ion sources (even for propulsion), and so on. A small mission of this character could very possibly accomplish more than an enormously more expensive and complex manned space station, over a period of a couple of years.

With the remote maintenance possibilities, one could expect very long functional lifetimes for such a craft.

2.7 Nuclear Power Industry

Although the very idea of master-slave manipulation originates in the environment of handling radioactive materials, and the present technology for hands is inherited mainly from the AEC-developed prototypes, this development has been handicapped by chronic funding limitations within the AEC itself. The first workable force-reflecting servomechanisms for master slave control seem to have come from this source, but even today their available manipulators are very clumsy and poorly articulated. The design of nuclear reactor components is still severely limited by the constraint of having to be serviced (if at all) by such devices, and it seems to be generally felt that the possibilities of intervention, in case of reactor malfunction, are dangerously restricted.

As for servicing small devices remotely, neither the AEC, NASA, NIH or ONR (who should be mentioned because of their prominent role in contracting for advanced remote manipulators in connection with undersea applications) have made any significant steps toward development of the smaller manipulators to which this proposal is directed. In particular, we wonder what plans exist for servicing small reactor components.

2.8 Agriculture

This is an area that might seem quite unrelated to the others, but it may be very important. If we look far ahead to an era in which small but capable programmed hand-eye devices are mass-produced inexpensively, we see applications that are very important. We mention this area with trepidation because it may seem too fantastic now to take seriously.

> To make the points plausible, we must interject a note on costs. Today, a research project involving a computer, mechanical hand and eye, and adequate staff, is a major investment, and one must get over the habit of thinking of such devices as multimillion dollar

> facilities. The cost of a small but powerful computer, in a decade, will approximate 1,000 dollars, and will be still decreasing. The cost of a high-resolution motorized eye will be less than 1000 dollars. The interface equipment will be the same order; the cost of a versatile mechanical hand with sensors is hard to estimate. But if they do find wide application in other industries, it is reasonable to expect the cost of a good hand also to come below 1,000 dollars, eventually. Therefore, it is reasonable to consider applications of hand-eye systems, before the 1990 period, at costs running well under 5,000 dollars, with complete mobility. This is less than a typical farm machine today, and much less than one with operator costs added in.

Such machines open the possibility of individualized, automated, attention to all phases of planting, soil preparation, inspection, weeding, fire fighting, de-insecting, harvesting, and new techniques for separation of plant parts. In some cases one could consider new methods of fertilization, unusual trimming of seedlings, new mass grafting methods, and application of small amounts of insecticides to specifically vulnerable anatomical sites on each plant. Large pests, such as caterpillars, could be removed individually without the use of any insecticides at all. We leave other facets of the concept of individual plant attention to imaginative agricultural specialists.

Another fantastic possibility may exist in the rehabilitation of acreage to new uses, by careful implanting of simplified artificial ecologies. In the near future, we can expect the emergence of enough mature agricultural science to design mini-ecologies for soil modification, in which a spaced mixture of a variety of different plant, insect, and bacterial ingredients can be expected to lead to a viable and useful territory. But the economic application of such ideas would require something like an inexpensive and versatile mobile hand-eye machine.

3.0 Foreign Competition

A nation can not expect to permanently maintain a competitive position, in the face of declining per capita productivity, by any combination of purely financial maneuvers, such as supporting industries whose balance is weak, or adjusting prices from above, or by manipulating tariffs and foreign exchange rates.

Our productivity has declined for several reasons. Our economists recognize social causes involving the needs and aspirations of Anerican workers. But they do not seem aware of the long-term effects of allowing basic industrial research to decline.

A chemical firm can benefit directly be investing in the development of a new product, technique, or even theory. A photographic firm can pull ahead by winning a new result from photographic research. But basic research in general production methodology only benefits everyone, not particularly the firm that pays the bill for learning how to do it. This means that no manager of a large corporation can rationally decide to make a large investment in basic industrial research! And no such manager has, in fact, done this.

In some other countries, there is a conscious awareness of this, and signs of systematic programs to do something about it. If other countries can advance, not merely in making particular products but in practical general knowledge about production itself, the effects on us could be very serious, since there is no reason to expect that the average American worker is going to become more dedicated to improving his own personal productivity!

4.0 Issues That Need Study

Here are some of the problems that must be solved to make this field develop more rapidly.

1. Which methods of machine vision will be the most effective?

2. What should be done to improve the mechanical dexterity of the manipulators?

3. What can be done to make the control programs more intelligent?

4. How can we improve education in Robotics at several universities?

5. Should there be an Institute for Research on Advanced Automation?

6. How can we develop laboratory equipment and systems to stimulate experimental work and training in this field?

7. What should be done about Micro-Automation?

This proposal is concerned primarily with the last two questions.

5.0 Current State of the Field of Robotics

Work on robotics is today at a natural turning point. Nearly enough is now known about computer-controlled manipulation to make it applicable in practical situations, but it will take a long time for this to happen unless the methods and the hardware are simplified, systematized, and made more widely available.

To describe the state of this art, it is convenient to summarize the situation in our own laboratory, which is probably the most advanced in at least some respects. A complete "hand-eye coordination" computer system has been constructed that can visually analyse a real-world scene composed of many simple arbitrarily-placed objects, and then physically rearrange that scene. For example, it can be told to make a copy of an existing physical structure using the same, or other parts. Because such a system must contain a great many sub-programs, each of which may be undependable and in any case depends in complicated ways on what the system "thinks" it perceives at each stage, we were not able to use conventional programming methods to obtain the required behavior. The new system is largely based on using the "pattern-directed" control methods of the new PLANNER language. It still has the status of a demonstration and research tool, with severe practical limitations, but nevertheless may be a turning point in learning how to put together subsystems that use very different kinds of information structures.

The practical limitations of this system are really very severe. The visual objects must be polyhedra -- objects with flat surfaces -- and optically very clean, because the visual analysis system does not know about textures or surface decorations. The manipulation abilities are equally rudimentary; the machine knows to grasp objects only between parallel plane surfaces, so cannot pick up most non-rectangular objects. These limitations apply also to the other robotics projects.

The language-understanding system of Terry Winograd is able to coordinate a substantial amount of knowledge about grammatical English with schematic knowledge about real-world physical structures, to produce the effect of understanding intentional and explanatory statements. Because the "meanings" of words and syntactic structures are described in terms of processes that interact both with the language-analysis and physical problem-solving aspects of the situation, the system can construct meaningful interpretations of statements, questions, and commands that might be ambiguous, incomplete, or ungrammatical by traditional linguistic criteria. This system, (or theory, as it more properly should be viewed) obtains a dramatic increase combinatorial search involved to test whether a given proposition is valid or not. It is intended to be a general formalism in which knowledge in a domain can be combined and integrated. Realistic problem solving programs will need vast amounts of knowledge. We consider all methods of solving problems to be legitimate. If a program should happen to already know the answer to the problem that it is asked to solve, then it is perfectly reasonable for the problem to be solved by table look-up. We should use the criterion that the problem solving power of a program should increase much faster than in direct proportion to the number of things that it is told. The important factors in judging a program are its power, elegance, generality, and efficiency.

in expressive power because of its use of "procedural definitions" in addition to traditional syntactic rules. By introducing a collection of "semantic specialist" program operators (whose activity is triggered by such events as occurrence of particular words during the "parsing" phase) we can handle a variety of problems connected with the semantics of a particular task environment. These processes convert the English text into new programs (in the PLANNER language) which obey commands or cause the execution of appropriate problem-solving or question-answering procedures. This makes it possible to solve linguistic problems that could not be handled by earlier syntactic theories. Much remains to be done in extending the semantic and reasoning resources of the system.

In our laboratory, the hand-eye and the linguistic control systems have not been mated together, mainly because we have been more concerned with extending each of them than with facing the problem of coordinating and standardizing two obviously inadequate subsystems. In the work at Stanford, such a coordination (of two much simpler systems) has been carried through on a demonstration basis, in fact, with the rudimentary English language semantic system preceded by a rudimentary speech-analysis system so that the robot can carry out certain kinds of very simple verbal commands.

The following laboratories have sophisticated robotics projects at the present time:

MIT Artificial Intelligence Laboratory (ARPA)
Stanford Artificial Intelligence Laboratory (ARPA)
Stanford Research Institute (ARPA)
Edinburgh Machine Intelligence Dept. (Scotland)
ElectroTechnical Lab., Tokyo Japan

There are also many small projects in universities and industrial laboratories, but these are all below "critical mass" and usually do not produce much. They start and stop. Some of the obstacles to rapid progress are:

Manpower --- Very few highly qualified people are working on advanced systems in this area. There are only a few places where one can learn about what is known, with access to appropriate equipment. The existing centers do not have enough staff to take on many more students.

Hardware ---The available optical and mechanical equipment is clumsy, unreliable, expensive, and requires special interfaces. To start a robotics project, today, means to begin an engineering project.

Software --- The programs and subroutines developed in the major projects are available only with difficulty, are undocumented, and run only on their parent computers. Higher level programs, mainly in LISP and related systems, are easier to transfer, but the lower level vision and real-time motor control systems depend on unique hardware and are non-transferrable. The most advanced new systems at MIT are in the new PLANNER language, which is becoming available slowly, but depends on large memory configurations that will not be generally available for a few years.

To get around these problems, we think the best strategy is to construct a self-contained research equipment package to make it easy to begin robotics work at other centers. It should include simple but adequate robot hardware, enough of a program library and local computation power for some research and applications, and provision for access to larger computation facilities via the ARPA network. The system we propose to construct will be inexpensive, relative to the present projects, and it will be small enough to fit in a large office, so that starting such a project will not be a major committment for an average university or industrial laboratory. By working through the network, we hope to avoid the problem of transferring high-level systems from one machine to another. We do not think the latter practical. The systems are changing too rapidly, depend on features of the local machine and time-sharing system, and the transfer could be made only by diverting the research staff involved.

6.0 What To Do Next?

We believe that the following two ideas are a basis for a good way to move ahead.

Thesis 1. Remote Operation is Feasible

We consider it practical to study advanced automation using a small computer that controls directly the robot hardware, does some local computation, and serves as a terminal for larger computations carried on elsewhere. Although real-time bandwidth problems limit some applications, many important applications are still feasible. High-speed real-time operation is not so important in this research as it might at first seem. For example, there is still a great deal of vision research that can be done using stored picture datam even in the case of moving objects.

Thesis 2. Small Scale Robotics has Important Advantages

We believe that the most practical physical size range, both for research and application, is the 0.001 inch resolution range. This scale is conveniently well-developed both in mechanical and optical respects. This decision will lead, we think, to neat and dependable equipment, and to highly motivated work pointed toward opening up new applications in the micro-scale areas. Some people have reservations about this, suggesting that in making the equipment small one may get involved with too many new problems at once. This is discussed in the next section. Certainly, a proposal to attempt robotics at a microscopic level would be open to this criticism, but we are talking about moving into a well established range of the low-precision machine tool industry, and not a hyper-modern microminiaturization area. The 0.001" size range has already so much established technology that it will make most mechanical and optical problems easier. The key points are that:

Precision machine tools: 0.001" is considered "low precision".

Optical microscopy: 0.001" is 50 times optical resolution.

This means that in both areas we remain far away from serious practical limitations.

7.0 What Size Range for the Mini-Robot?

All current robotics projects are scaled to a working area of the order of one meter, with mechanical resolutions (open-cycle) of a few millimeters. This is also the general range of the large, standard, industrial manipulators of the Versatran and Unimate class. Of course, it is also the scale that humans like, when working without the use of instruments.

We will not discuss larger size ranges, such as might be useful in construction, warehousing, etc. But let us review the problems of various smaller scales.

One millimeter resolution

It would be very difficult to make the present projects' equipment work in this or any smaller range. The mechanical design of the commercially available mechanical arms cannot be refined to this degree,because of problems with structural deformation and stability of the positioning servos. Optically, there would be no problem.

One-tenth millimeter resolution

The mechanical problems with present equipment would be even worse. If the equipment were to be redesigned, this size range would still be illogical, because there would be no problem in reaching the next size range discussed below.

One-thousandth inch resolution (.025 millimeter)

0.001" is considered "low precision" in the machine tool area. The entire spectrum of equipment from the machine tool industry is available, from milling machines to measuring devices; effectors and sensors. It is easy to make new devices with this tolerance using ordinary machine tools. (It is not much more difficult to reach tolerances of 0.0002 or so, with higher precision tools, to provide the margin one needs for making complicated 0.001 devices; it quickly becomes much harder to reach higher precision than that.) One possible product of this whole program, of course, would be microminiature machine tools of still higher precision.

Optically, things are also extremely favorable for this size range. While it is theoretically possible to resolve spatial intervals smaller than 0.000040" (0.001 mm) in visible light, one wants to say away from fundamental physical limits by a sizable factor. To reach limits of optical resolution one needs special equipment, oil immersion, careful control of lighting, and extremely inconvenient restrictions on depth of field and clearance. The 0.001" resolution is 50 wavelengths, and can be reached by simple, inexpensive, magnifying devices with large fields of view and arbitrary working distances

The 0.0001" range.

Here we would be in serious problems, both optically and mechanically. 0.0001 is only four wavelengths of yellow light, and would be a very high machine-tool precision. Sensors would be rather hard to obtain. This range is not out of the question, because the technology of biological micromanipulators operates there, and we could extend those devices. However, it is hard to see any technical advantages to doing that in this early stage.

0.00001 and below.

By using such special devices as the scanning electron microscope, and pantographic mechanical devices, one could reach further. But the technical problems would be enormous, today. These will be practical domains, however, after a few years of experience with the 0.001 field.

Conclusion

The most effective configuration of equipment, both for developing automation and for reaching toward micro-technology, is to attack the 0.001" size range using as much as possible from today's methods in industry, machine-tool, microscopy, and sensors from biology and medicine.

The most straightforward configuration is to base the mechanical system on a six-axis motorized arrangement of machine-tool slides and rotary components. On this, a small mechanical hand must be mounted. The latter has to be developed. In any case, it should not be difficult to bring such a system up to the performance of the parallel-jaw grasping systems currently in use in the AI projects. Study contracts should be started for deciding on the next step. Study of micro-sensors should also begin soon.

It will take some time to consolidate robotics knowledge and practice for any application whatever, and the concurrent development in the 0.001" size range need not become a diversion, since the problems are substantially independent.

The next step, of pushing into the ranges of 0.0001 and below, will require careful planning, and probably would be done most profitably with a particular application in mind. This is because more of the techniques will be new and specialized, and one would need a clearly defined particular application to motivate decisions about how to proceed. Our first candidate for such a goal is the assembly of computer hardware from semiconductors, with the micro-hand-eye system doing all the mounting, positioning, wiring, and physical assembly.

8.0 Modules for the System

To free research workers from the inadequacies of one-of-a-kind systems, we want to develop and make available certain modules.

> The modules should be easy to assemble from available components, with very little special engineering.
>
> The modules should be small enough that a significant robotics project can exist in an ordinary room, In fact, our goal is to make the whole system desk-sized.
>
> A complete system should not be extremely expensive. We would like a copy to cost less than 100K. Incidentally, this goal is parallel to another A.I. Laboratory goal (under NSF support) to make a prototype multi-console system for schools, which also controls a substantial amount of motorized real-time activity in local devices. In fact, the "controller" discussed below is being developed jointly with that project.
>
> The system should be fully capable of operating, with reduced bandwidth, from remote computation services

over systems like the ARPA network.

The modest physical size and simplicity of the system should help to reduce the general overhead of work in this area, and make it possible for small university and industrial laboratories to participate.

8.1 Module 1. The CONTROLLER Device.

The Controller is a versatile interface device designed for controlling peripheral devices via direct or dataphone lines to a computer. Basically, it is a character-string device. Certain characters have control mode effects on the controller while the remaining characters are treated as data. The controller has a number of input and output "ports" and, depending on the mode the controller is in, data is serially multiplexed between an arbitrary subset of currently active ports.

Each port can be equipped with any or all of these features:

Digital data output with strobe pulse.

Any number of ports can use the same strobe pulse, if one wants "multiple precision" data.

Analog input and output. We plan a choice between ASCII-based low precision and mini-computer word based high precision data.

Direct control of digital stepping motors. The controller has a microcode for fast positioning of external devices.

The controller is fast enough to multiplex data at any available line bandwidth. Its microcode is efficient in that a single character can interrupt the multiplexing mode and transmit a data stream to any selected port at full line speed, until interrupted and returned to multiplexing mode by another special character.

For input one can select a cycle of input ports to be scanned at a regular rate. Alternatively, one can use a simple priority system so that an external device that wants attention can seize a port and transmit a data stream at full rate until it is through or a higher priority device seizes another port.

A prototype of the output part of the controller already exists. We used it in July 1971 to give a demonstration in Berkeley, California, in which our PDP-10 timesharing computer in Cambridge was programmed to operate simultaneously a keyboard-printer, two stepping motors and four other output channels. The demonstration involved a mechanical "turtle" moving and dancing to some programmed music, to show educators some of the possibilities of using unconventional output devices for programs written by young children. We had no difficulty operating these devices simultaneously, at very impressive speeds, over an ordinary phone line. Although the character rate was just 30 characters/second, the microcode operates the stepping motors at eight times this rate.

We propose to complete and package the design of the Controller. The whole device will be quite inexpensive, probably under 5K final cost, including the telephone interfaces and the stepping motor power and controls. Among other features, the controller has a set of relay-switched ports and indicators, which can control other devices, including A.C. power.

The Controller is designed to operate either as the I-O interface for a minicomputer, or by itself using only a telephone connection to a time-sharing system. For the mini-robot laboratory, the controller should have no difficulty handling the hand-eye complex, the console, and a few other devices. The current version of the controller handles only 7-bit ASCII words; we will have to make a 16-bit version of the controller for this project.

8.2 Module 2. The Mini-Manipulator System.

Even on the macro-scale, manipulator hardware is in poor shape today. There is a real need for development of a much better arm-hand system with good force and pressure feedback actuators and a relatively unobstructed working field.

For this project, we plan to produce a simply articulated hand-arm, with good force-position actuators and sensors. The design will depend partly on the results of a semi-public design competition. Following this, we will initiate planning studies of a much more dextrous hand design, in cooperation with the other robotics projects.

We plan also to work toward an easily applied laboratory "kit" of motion control devices, for making mechanical experiments easy on the non-miniature scale. The A.I. Lab plans to develop such a "kit", eventually, in connection with its elementaryschool hardware accessory project. These will perhaps serve as first-step prototypes; although they will not meet industrial specifications, the requirement of being child-proof is pretty strong.

While the design of the manipulator is not yet decided, the most straightforward system would seem to be one in which a simple hand is supported by a six degree-of-freedom base. This could be a commercial three-axis rectangular drive, of which there are many available, plus three smaller rotator modules, which are not so easy to obtain. We plan to drive these open-loop, without force feedback. But this base provides a hand mount that can be positioned and oriented freely, so at the mounting base we would provide a sensitive six-axis strain-sensor with which force control can be realized.

The hand itself would attach to a thin rod, and have a simple hemispherical-mobility wrist, actuated by tendons from the mounting base. Thus there would be no complicated actuators within the miniature part of the system. Finally the hand would have a simple grasping action, tendon actuated also.

We will also consider other designs in which the manipulative mobility is more divided. One might split the mobility, with some degrees of freedom operating a work stage, while others move the hand. There may be advantages in splitting the system into two three-axis devices instead of one six-axis device. This would improve rigidity and simplify mechanical and information transmission problems. It might also make it practical to place more hands and eyes in the same field, assigning only three motion degrees to each. Disadvantages include the effects of motion on the operating "world" and the non-independence of the several devices.

Sensors on hands

There are many possible approaches to instrumenting small mechanical hands. Some of the possibilities are:

1. Standard devices, such as miniature strain guages. In the medical area, there exist very small probe sensors for use in cardiac catheters, etc., to monitor pressure and temperature. The basic semiconductor sensors are already quite small.

2. Fibre-optic methods. A single light fibre, or a pair, can be used to transmit information from any sort of mechanical detector by occluding the back-reflection or transmission into another fibre. This is quite attractive becausel it is very easy to do the instrumentation at the computer end of

the fibre(s), where there is plenty of room, even if the diameter of the fibre is very small.

3. As we noted in discussing spaceship applications of robotics, there are new possibilities for using sensors that are passive in the sense that they need not transmit data back along any direct channel. Instead, their state could be observed by the computer eye. While this idea is probably inappropriate for our moderate size-scale hand system, it is something to explore for future microminiature applications. The space position of a simple retroreflective glass bead can be measured to 1 micron by now available interference methods.

8.3 Module 3. The Computer Eye.

Suitable "computer eyes" are still not available at reasonable costs. A series of industrial eyes would run over a wide range of capabilities. The simplest applications could use small arrays of separate detectors, with a multiplexing device. A deluxe eye needs high precision and resolution, but this could be a joint outcome of optics and mechanical devices. Developments in solid state photosensitive arrays, or new image tubes, may turn out to be the best thing, and we cannot decide yet.

The eye we plan to use in this project is simpler than those we have used in the past, and a few remarks about this are appropriate. The orientation of work on computer vision, in the ARPA projects, was never intended to be directly applicable to industrial problems. Many of the serious problems encountered in those experimental research systems can be sidestepped by using illumination and scanning schemes that yield three-dimensional information directly. In this way, high performance can be obtained by programs that we expect will be much easier to develop and debug. In the Artificial Intelligence projects, it should be understood, the goal of the vision experiments was to face certain very perplexing problems, such as --- "how can one use his knowledge about the world to reconstruct a three-dimensional scene from a single two-dimensional picture." The resulting computer programs were indeed able, in some cases, to deduce the spatial position of objects, using elaborate chains of reasoning about which objects appear to support which others, where an object probably rests on the table, and even using information about what object could have caused such-and-such a shadow. We could have avoided this difficult kind of question by using direct three-dimensional range-finding method, such as the one used by Horn and described in our 1969 progress report. In the industrial situation one wants to get the best information one can as directly as possible. One might consider:

Stereo
Focus measurements.
Interference methods
Controlled scanning beam

The latter was used by Shirai, of the ETL in Tokyo, who is a visiting scientist at our laboratory for 1971-1972. Such schemes are practical, yield great precision, simplify the photodetector problems in most cases. Of course, they tend to be limited to close range applications, but that is what we want now.

At this time, the most plausible design for a computer eye for automation applications would seem to be one in which the light source projects a slit of light in a program selected vertical plane, while the receptor measures where the light appears in each horizontal plane. If the source is mounted oblique to the receptor, the horizontal position of the intersection yields a direct measurement of the spatial location of the illuminated point, by simple triangulation. It is hard to say what is the best engineering realization of this, but one scheme that seems practical would be a CRT or laser light source and a vidicon photodetector. We have not selected the slit deflection system yet.

The photodetector would be equipped with a special interface in which a peak detector finds the location of the intensity maximum along each horizontal sweep and records the position and intensity in a special memory. Thus, each sweep gives a vertical cross-section of the spatial situation. A small number of such scans would be adequate for many positioning and inspection applications, and would take only a few TV scans. This system alone should handle many problems. We hope to add to it a conventional image acquisition option, using the same sensor.

In any case, the decision will depend also on the outcome of the sensor study being carried out now at Stanford.

9.0 Programming

Standardizing enough basic vision and manipulation capabilities is a major research problem. Adapting "general" methods to particular jobs involves problems more like those in the "children's story" and "common-sense" areas of Artificial intelligence research than like those in the automatic machine-tool or programming language compiler areas. Physical situations have to be described to the machine in terms of goals and purposes, as well as in terms of shapes and shadows. If the programming job for each application is not to be impossibly lengthy and expensive, such programming systems will have to have very smoothly debugged collections of real-life knowledge, so that they do not have to be told huge amounts of trivial advice. Soft surfaces should be handled with soft effectors; things slide down highly tilted surfaces (and balls roll off slightly tilted surfaces); one should not close a box until it has been filled, etc. This is all obviously a matter of degree, and in a particular situation few such problems may arise. But in every real situation, one finds some such problems.

This proposal does not cover doing much of this. Some will be done in the main AI project, but our real goal is to make it easier for other groups to participate in this open-ended area.

The system minicomputer will certainly need to have good motor-coordination programs capable of grasping objects and using available sensors. The visual programs must be capable of elementary three-dimensional analysis of suitable scenes. These programs must have interfaces that make them accessible to higher-level programs in the remote computation centers, using descriptions concise enough to avoid fatal bandwith limitations.

10 Costs

VISION

For practical applications it is important for the visual system to be capable of reasonably direct rangefinding, and we expect to use the programmable slit scheme described above. Further study is needed to decide on the best system, and it is hard to estimate the cost, but the programmable slit light source, a good CCTV eye, and the video-processing interface should cost 50K to develop and less to duplicate.

OPTICS

Motorized optics should cost about 6K. It may be useful to build a flexible remote-controlled head, perhaps using fibre optics, for vision in tight spaces, but that can wait.

MECHANICAL

The basic six-axis mechanical base should cost about 30K to make or buy.

We plan to study available micromanipulators, as candidates for holding the hand --- 5K acquisition cost.

We will test a number of special surgical instruments as parts of end organs --- 1K.

The simplest hand that makes sense would have a parallel jaw gripper and a spherical-mobility wrist close to the gripper. It might make sense to investigate a continuously-deformable structure for the wrist, easy to fabricate and potentially extensible to the micro-dimensional region. Engineers at Draper are interested in this design problem.

The axis to the wrist should be instrumented with the six-axis force-sensor.

The first prototype hand will cost about 25K to develop, including motorized wrist and gripper. We will begin with an earlier, larger prototype. Later, we would want to have the hand redesigned by an outside firm, and built for reliability in as small a size as is feasible without running into serious problems. We are considering engaging the talent of Ralph

Mosher, who built many of the most advanced manipulators of the past, when he was at General Electric. He has also expressed interest in this job.

We include in the proposed budget 25K for the various design competitions, to allow possible subcontractors to submit well worked-out proposals.

When the hand system is available, it will find many applications as a simple master-slave device, without the computer. We will want to have a master input harness for the human operator in any case, and this may cost perhaps 5K to build, or to buy, if an available master control from a different arm is suitable. At a later stage, if the mini-hand turns out to be generally useful, a bilateral force-reflecting servosystem should be developed for it, but not as part of this contract. The design and human engineering of this controller is difficult. Again, a design competition would be appropriate, and funding sought from the appropriate application agency, NASA, AEC, or NIH. The problem of sensory instrumentation of the hand is not well understood. We hope that a six-axis force sensor at the wrist will provide enough feedback for many applications -- 20K to develop or purchase. It would be interesting to explore the fibre-optic "sensory nerve" idea, but this might be more appropriately subcontracted to another laboratory; B. Howland at Lincoln is interested.

CONTROLLER

The controller development should cost about 25K to carry it from its present state to a version with adequate motor control and input multiplexing ability.

COMPUTER

The minicomputer for the system, to make low-level vision programs, etc., available to a Network type of operation, will cost about 50K, mainly for memory. This does not include cost of providing a suitable local display, which we estimate at another 15-20K.

PROGRAMMING

The programming costs would be large in an industrial situation, but much of it can be done here by students, supervised by the laboratory staff.

The project needs something like 3 engineer/designers, a systems programmer, a manager, several students and assistants. Some of them will be shared with the AI project, but the cost should be about 90-100 man months. We are estimating 1.5K cost per man-month direct salary plus benefits and overhead. This low value is justified by the fact that the A.I. lab has been able to attract highly capable students to do advanced work, and we are counting on the attractiveness of this project to make this budget work.

GENERAL COMPUTER SUPPORT

The additional load on our central computer facilities will require some support in consoles, and primary and/or secondary memory. It is difficult to pinpoint this because it will depend on where the extra traffic has the worst effect. We are at present extremely weak in the graphic console area, and in primary memory. The project will be the main job of our existing shop facility. The shop facility will also be shared with the NSF contract that we expect to renew, for developing the mini-computer system for the LOGO education project, and the usual maintenance of the other AI lab projects. These general costs approximate 100K.

In the estimates below, the first column is estimated cost of hardware and materials. The second column is Man-Months of engineering work, and the third column is Man-Months of programming for the system. The proposal does not cover the cost of converting the low-level vision programs and the required new programs for the mini-computer system; these will presumably happen as the system is completed, by students associated with the Vision group.

VISION

------------------ good TV camera

5K

------------------ engineering for stability and corrections

5K 3MM

------------------ video processor and interface

10K 6MM 2PP

Programmable light source

------------------ Laser and deflector

8K 2MM

------------------ interface for light source

2K 2MM

OPTICS

------------------ Microscope parts

5K 3MM

------------------ motorized controls (4)

6K 2MM

Manipulator

------------------ Micromanipulator study

5K 1MM

------------------ Two 3-axis machine tables

10K

------------------ motorized controls (6)

15K 2MM

HAND 1

------------------ prototype of flexible hand

2K 3MM

------------------ simple 3 axis wrist

2K 3MM

------------------ motorized controls (6)

9K 2MM

------------------ master slave control

5K 2MM

------------------ surgical end-instruments

1K 1MM

Hand 2

------------------ Design competition

25k

------------------ Outside contract for mini-hand

25K

SENSORS

------------------ 6-axis wrist sensor

20K 3MM 3PP

------------------ Grasp Sensor

2K 1MM

CONTROLLER

------------------ minicomputer with memory

50K 2MM 18PP

------------------ interface to controller

2K 1MM

------------------ console

15K 1MM

------------------ direct interface to arpa network

2K 1MM 2PP

------------------ input multiplexor

5K 6MM 2PP

------------------ output multiplexor

5K 3MM

------------------ final prototype packaging

10K 3MM

------------------ general purpose mechanical modules for robotics research

10K 5MM

General laboratory and computer support 100K 6PP

includes additional memory, consoles, shop support

construction of mini-robot final assembly

------------------ Management

20K

------------------ Rent, etc.

25K (see budget below)