

**A LABORATORY ENVIRONMENT FOR
APPLICATIONS ORIENTED VISION AND MANIPULATION**

by
Berthold K. P. Horn
and
Patrick H. Winston

Abstract

This report is a brief summary guide to work done in the M. I. T. Artificial Intelligence Laboratory directed at the production of tools for productivity technology research. For detailed coverage of the work, readers should use this summary as an introduction to the reports and papers listed in the bibliography.

Support for the laboratory's research described here was provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research Contract number N00014-70-A-0362-0005.

TABLE OF CONTENTS

APPLICATIONS ORIENTED VISION AND MANIPULATION.....	3
MILESTONES	11
RECOMMENDATIONS FOR THE FUTURE	17
TECHNOLOGY TRANSFER.....	19
THE REPORTS	21
BIBLIOGRAPHY OF MEMOS & TECHNICAL REPORTS	23
WORKING PAPER BIBLIOGRAPHY	27
APPENDICES	29

This report is a brief summary guide to work done in the M. I. T. Artificial Intelligence Laboratory directed at the production of tools for productivity technology research. For detailed coverage of the work, readers should use this summary as an introduction to the reports and papers listed in the bibliography. Other relevant material is in our general progress report, *New Progress in Artificial Intelligence* (AI-TR-310), and in our proposal for work in 1976 from which this summary was partly adapted (AIM-366).

APPLICATIONS ORIENTED VISION AND MANIPULATION

In our proposal (AIM-274), we emphasized two points:

- We argued a need for a workable, inexpensive set of hardware and software tools for doing research in applications oriented computer vision and manipulation. We laid out a plan for the creation of these tools.

- We suggested that tool creation without one or two target domains can be inefficient and leads to poor resource allocation. We specified the problem of electronic circuit card manufacture and repair as suitable for bringing general issues into crisp, manageable form.

For the most part we are satisfied with the progress in these areas. The inexpensive, exportable vision and manipulation laboratory now exists. It has demonstrated its usefulness by supporting important target domain projects involving electronic circuits and by facilitating first steps outside electronics in other worlds of interest like the world of small mechanical devices.

We now give a brief description of the general plan of which this work was a major part, proceed with an itemization of the laboratory's components, move then to a brief discussion of some progress made with it, and then finish with some recommendations for future research.

Productivity Technology Has An Important Mission

Broadly speaking, the overall goal of our work in applied vision and manipulation has been to develop uses for artificial intelligence techniques in industry. We have had this goal for the following reasons:

- Maintenance of the economic strength of the country in the competitive international arena through automated production, particularly in assembly.
- Creation of a more flexible brand of manufacturing capable of very quick response in times of rapidly changing need.
- Preparation for automated mining, both underground and undersea, as well as automated farming, space exploration and recovery, ordinary maintenance, and maintenance in hazardous environments.

We believe these objectives are vitally important because preeminent computer science and industry is one of the major strengths of the United States. We believe we must translate this advantage into corresponding capabilities in areas where our world position is eroding, especially in productivity. We think changes will come about quickly when a certain threshold of technology is passed over. Just as pocket calculators swept the country like a tidal wave, we believe computer vision and manipulation will soon sweep through the manufacturing and maintenance industries. But before discussing what has been accomplished, let us look at the overall path and see how it reaches toward the stated objectives.

Progress In Transferring Artificial Intelligence Ideas To Industrial Practice Involves Four Major Steps

Our work has been driven by the following plan elements:

- 1 Pick the simplest possible domain, do vision and manipulation in that domain, and use the results to measure the difficulty of the problems in order to allocate resources and predict when applied results should emerge.

This had already been accomplished prior to our direct efforts on applications. The blocks world was the first domain; the difficulty of vision dictated the concentration of effort on it and the relative neglect of manipulation. It became clear that the existence of global knowledge and the non-linearity of noise effects render most previous work on image processing of little value. Instead, problem solving, constraint exploitation, and symbolic low level processing became central issues. It also became clear that vision is very hard, and for the sake of near term application results, the goals in vision research should be divided into no-tricks basic study and no-holds-barred, special-purpose, domain-dependent work.

- 2 Once the problems are uncovered through a struggle with the "simple" domain, then it is time to get together hardware and software tools capable of meaningful, cost conscious prototype arrangement to solve particular problems.

This too has been accomplished and stands as one of the major products of the research reported here. Our Micro-Automation Laboratory exists and supports applications work. It is described in a moment.

- 3 Use the tools to tackle real problems. At first these may be selected because they are amenable to solution through the existing tools and techniques -- obvious indisputable commercial viability is a secondary objective.

Our work on electronic assembly was chosen for suitability in this phase. Many results, described later, are in hand.

- 4 Finally, apply the technology on the most serious problems where the saving would be greatest in human and fiscal terms.

This is the next step, one for which our work has prepared us. More specifically, this translates to the assembly of mechanical devices. The tough assembly problems involved seem essential to industry and appropriate for solution by some combination of vision, force sensing manipulation, and

now-understood Artificial Intelligence programming ideas.


The Hardware Is Defined

The laboratory for productivity technology is built around a PDP11/40 with extended instruction set (EIS), 31K of core memory, two 1.2 megaword disk drives, and a GT40 graphics display terminal. In addition, we have the following devices which we believe form a good checklist for getting started in productivity technology research:

- Two VICARM mechanical arms, 3/4 human size, with six joints and six degrees of freedom, position and velocity sensors in most joints, and gripper hands. These comprise our primary manipulation system. Texas Instruments, the University of Illinois, and the Naval Research Laboratory have ordered the MIT designed arm and more are expected to follow suit. Figure 1 shows the general appearance of these arms.

- A Vidicon picture digitizer, consisting of a television camera and a television monitor, and a digitizer which converts the scanned picture to 8 bit digital data at a maximum rate of 65 microseconds for each point in the TV scan. This is currently our primary vision system. (Soon solid state arrays will be the better choice, but not as of this writing.)

- A Hughes modulatable laser with an output of .001 watt, used principally as a pointer or tracking device.



1/4 scale

Designed at M.I.T. by Vic Scheinman

- Two motor-driven mirror-deflection systems, by Spatial Data Systems, Inc. Each mirror system has two independently controllable mirrors with perpendicular axes of rotation, and settling times of approximately 3 milliseconds for motion of the mirrors. One system is used to deflect the Hughes laser; the other is attached to a PIN diode via an optical system and is a slow but very linear, high resolution vision input device. *The inferiority of even good vidicons is immediately obvious from comparisons of vidicon pictures against those made with this mirror scanner.*

- An Optronics International, Inc. photowriter, capable of making extremely high resolution photographic negatives on film from stored digital pictures. *We believe having such a device is essential for research in image analysis and understanding.*

- A precision X-Y table, made by the Icon Company, to our specifications, capable of motions in two orthogonal directions with an accuracy of .001 inch, over a range of motion of 6 inches with a maximum velocity of 4 inches/second. This device is used to position work under the mechanical arm or the vision system.

- An Analogic digital-to-analog and analog-to-digital converter with 8 D to A channels and 64 A to D channels, capable of performing data conversions in 3 microseconds, and to which we have added an additional 16 D to A channels.

The Software Is Done

The Productivity Technology Group PDP11/40 uses standard DEC software, the Disk Operating System (DOS), in an effort to make software developed within the group useable elsewhere. The group has developed the following software:

- VIDIN is the visual input program. It is used in conjunction with either the Vidicon or the PIN-diode/mirror system to obtain digitized pictures which are stored on disk. The user can select arbitrary sub-sections of the entire picture being scanned, and the program continually displays to him the sub-section of the picture which he has selected.

- MAPPER is a program for computing statistics on stored digitized pictures. MAPPER can provide histograms and intensity gradients, printouts of the raw data and even graphic representations of the picture itself. Operating in conjunction with a program that runs on the GT40 graphic display terminal, it can threshold a picture into 8 intensity gradations, and then cause the picture to be displayed on the GT40. MAPPER enables the user to examine small subsections of the entire digitized picture for detail, or to combine subsections together to get a larger view of the statistical analysis of the picture.

- LISP, a small, multiprocessing LISP interpreter, with shallow bindings and a non-recursive garbage collector which conserves memory. This LISP does have the restriction that atom print-names must be 3 characters long, or less. It does standard DOS file I/O, and can additionally do input or output from a variety of terminal devices, or from the Artificial Intelligence Lab's PDP10 timesharing system.

- RUG, a symbolic, DOS-oriented debugger, which replaces the DOS debugger ODT. Using RUG, a user may debug standard DOS format programs; the debugger makes available to him all his symbols, disassembles the contents of core into assembly language statements with symbolic references included, and allows him to enter assembly language code into core by merely typing it in. RUG has various numeric and special-purpose type-in and type-out routines (such as ASCII and RADIX-50). It allows the user to set breakpoints in his program, and additionally allows him to monitor the contents of a core location, register, or device for some desired condition.


- A set of arm control routines for the VICARM mechanical arm. These are implemented in LISP, and include routines for moving the arm to a particular state-vector (essentially a position and velocity for each joint), and for halting the arm's motion and holding it still. The routines incorporate a trajectory planner which decides on the route the arm will take to get to the desired state-vector, and a dynamic routine. The dynamic routine looks at the desired state-vector as computed by the trajectory planner at the end of the next time increment, and computes the torques which must be applied to each joint in order to achieve that state-vector.

Finally, there is an improved editor which utilizes the graphic capabilities of the GT40 terminal to implement a real-time edit mode, and there are programs for transmitting files to and from the timesharing system, for displaying images on our GT40 terminal, and for other typical system activities.

MILESTONES

The flow of basic Artificial Intelligence techniques to applied problems has begun in the Artificial Intelligence laboratories and a complete list of achievements in this direction by the community would surely include the impressive work done in Japan on A.I. inspired ZIP code readers and circuit card inspectors, the work at Stanford, Stanford Research Institute, and the Draper Laboratory on manipulators and manipulator control, and our own work, of which the following are of particular relevance:

- Inoue and Silver demonstrated the automatic assembly of a radial bearing with 12 micrometer tolerances, thereby demonstrating that *complicated assemblies can be made using flexible, programmed manipulators*. The achievement further demonstrated that *knowledge based force sensitive manipulators are the key to advanced assembly operations*, and that *high level programming languages make programming for mechanical assembly a viable concept*. The bearing is illustrated in figure 2.
- Horn showed how to apply the technique of *knowledge based profile analysis*, borrowed from earlier blocks world work by Binford, Shirai and others, to the applied problem of orienting integrated circuit chips in preparation for standard lead bonding. Figure 3 shows some typical profiles encountered.
- Lozano also used knowledge-based profile analysis to locate the resistors and capacitors on a circuit board in preparation for future work on automatic testing and repair.
- Mitnick found a way of inspecting integrated circuit lead frames for bent finger defects. Lead frames are small stampings which are sometimes used to connect oriented chips to the D.I.P. contacts.



Size and clearance of the parts.




Illustration of the waveform and measurements taken to determine a hypothesized edge position.




Illustration of the waveform and some of the measurements taken to verify a hypothesized edge position.


- Taenzer used *region growing ideas* to inspect solder joints for the standard defects -- pin holes, cold joints, and open holes. See figure 4.

So far the programs of Mitnick and Taenzer are too slow for immediate commercial use -- their point was to demonstrate feasibility.


Other work now being performed using the laboratory's assembled hardware and software tools include the following items which further illustrate the flexibility achieved.

- Speckert has incorporated *higher level knowledge* about obstacles in low level tracking routines that follow objects through otherwise difficult accelerations. The mirror scanner system is used for this purpose. Figure 5 shows a trace of Speckert's points of observation as his system tracks a ping-pong ball through a bounce. Other tests of the tracking techniques include experiments in which walking people are tracked in order to help understand biped locomotion.
- Woodham is using the vidicon-based vision equipment in studies aimed at automatic inspection of modern jet engine castings.
- Horn is using the image output device in studying techniques for understanding shading such as is required for doing automatic cartography.
- Raibert is using the manipulation equipment to study new methods for manipulator control based on automatic learning about arm dynamics through ordered experiments in state space.

Solder joint with pin hole



BALL NEARLY STOPPED




LOCUS OF POINTS SCANNED WHILE VISUALLY TRACKING
A BOUNCING PING PONG BALL

RECOMMENDATIONS FOR THE FUTURE

It seems sensible to us to implement a system which builds a small gasoline engine -- very likely the model airplane engine illustrated in figure 6 -- with a view toward solving the following problems:

- The programming problem: Our intention is that the assembly program should look like an annotated parts list with advice about the assembly, not low level instructions. We are inspired in this effort by complementary work at Stanford and by our own history of success with very high level languages like PLANNER and CONNIVER.
- The parts orienting problem: The parts are to be scattered at random on the table or drawn from bins. The SRI results will be adopted, if sufficient.
- The flexible part problem: The engine has a gasket. We hope to connect up the fuel line as well. We believe our handling of flexible wires in our electronics assembly work will be valuable experience.
- The two-hand coordination problem: The engine has a difficult problem inasmuch as piston, piston pin, and connecting rod must be brought together. Inoue's and Silver's work with the radial bearing, presented in the progress report, gives us some solid close-tolerance manipulation experience.



- **The straight coordinated movement problem:** In an effort to eliminate the complex equations normally required for dealing with arm dynamics, Raibert is constructing a model whereby the arm controller learns from previous attempts at movement control and stores this learned behavior into a large data base on disk. The first step is to cause slow motion of the arm between pairs of points in order to teach the system the gravitational parameters along the path in between. Then the arm controller can begin to learn about the dynamics of motion along the same paths by iterating the motions at speed, trying to learn about the arms dynamics in order to reduce the gap between the previous iteration and the desired motion.

- **Coordinated vision and manipulation:** Insertion of the crankshaft into the connecting rod offers this challenge. The try will not win every time, we suspect, bringing in the opportunity to deal with fault detection and recovery in a natural way. Speckert's tracking work may be useful here, as well as Taenzer's inspection and insertion work.

- **Tools:** The entire job will require some serious attention to thought about the manipulator's tools.

- **Inspection:** We may, if fortunate, be able to test the assembled engine and even adjust the needle valve. Inspection for surface defects, however, is to be done in the casting domain, described later.

Within constraints imposed by limited resources, we intend to continue work in these areas, with the basic supporting tools now seeming adequate for the task.

TECHNOLOGY TRANSFER

Of course, industrial exposure has been vital to this program because sensible resource allocation and technology transfer are so important. MIT's very active Industrial Liason Office has been very helpful in this

regard by helping arrange for us a wide spectrum of contacts in industry. Production facilities visited include those of Texas Instruments, Hewlett-Packard, Data General, Digital Equipment, IBM, the RCA hybrid circuit facility, the Kennecott Copper Company Peabody coal mine, General Motors car assembly plants, Delco integrated circuit facilities, AMP incorporated electrical connection and electrical assembly facilities, the Foxboro industrial controls facility, United Shoe Company's automated electronic equipment machinery production, and the GCA facility charged with the production of advanced lead bonding equipment.

In addition we have engaged in several steps which we believe considerably shorten the time it would otherwise take to bring vision and manipulation work into practice on the factory floor. We recommend similar steps for others engaged in productivity technology research.

- We have introduced a new MIT course on the subject of computer vision and manipulation in order to prepare new graduates for active leadership in the field. A text has been undertaken as a mechanism for encouraging such courses elsewhere.

- We have introduced a summer session course for the MIT industry oriented summer program. There was a favorable first and second year turnout with engineers and managers from Boeing, Honeywell, New York Telephone, The U. S. Department of Transportation, The Bureau of Mines, The Office of Naval Research, American Can, Modicon, Allied Technology, Western Electric, Westinghouse, Kodak, and Xerox attending. General Motors has purchased our Videotape introductory Artificial Intelligence course.

- We have whenever possible built the vision and manipulation laboratory out of easily obtainable equipment and we have avoided modifying this equipment. When new equipment has had to be designed, we have avoided one-of-a-kind in-house construction.

THE REPORTS

Most of the papers enumerated in this section were produced as a direct consequence of our work in productivity technology. Others were written in connection with other basic research foci of the laboratory which treat topics of interest to readers concerned about design automation as well as production *per se*.

Five representative papers from the list are reproduced in the appendix. They cover part of the work on manipulator geometry, manipulator control, force feedback assembly, visual inspection, and knowledge based understanding of engineering design facts.

Laboratory memoranda relevant to productivity technology:

267 Manipulator Design Vignettes, Marvin Minsky, October 1972. (\$1.30)

This memo is about mechanical arms. The literature on robotics seems to be deficient in such discussions, perhaps because not enough sharp theoretical problems have been formulated to attract interest. I'm sure many of these matters have been discussed in other literatures -- prosthetics, orthopedics, mechanical engineering, etc., and references to such discussions would be welcome. We raise these issues in the context of designing the mini-robot system in the A.I. Laboratory in 1972-1973. But we would like to attract the interest of the general heuristic programming community to such questions.

273 The Little Robot System, David Silver, January 1973. (\$.90) AD-773-929.

The Little Robot System provides for the I.T.S. user a medium size four degree of freedom six axis robot which is controlled by the PDP-6 computer through the programming language LISP. The robot includes eight force feedback channels which when interpreted by the PDP-6 are read by LISP as the signed force applied to the end of the fingers.

c289 Visual Position Extraction Using Stereo Eye Systems With A Relative Rotational Motion Capability, Daniel W. Corwin, March 1973. (Originally issued as Vision Flash 22, January 1972.)

295 On Lightness, Berthold K.P. Horn, October 1973 (see Computer Graphics and Information Processing, December, 1974). (\$2.10) (AD-773-569).

The intensity at a point in an image is the product of the reflectance at the corresponding object point and the intensity of illumination at that point. We are able to perceive lightness, a quantity closely correlated with reflectance. How then do we eliminate the component due to illumination from the image on our retina? The two components of image intensity differ in their spatial distribution. A method is presented here which takes advantage of this to compute lightness from image intensity in a layered, parallel fashion.

300 Design Outline for Mini-Arms Based on Manipulator Technology, Carl R. Flatau, May 1973 (issued as A.I. Memo January 1974). (\$1.70) AD-773-570.

The design of small manipulators is an art requiring proficiency in diverse disciplines. This paper documents some of the general ideas illustrated by a particular design for an arm roughly one quarter human size. The material is divided into the following sections:

- A. General design constraints.
- B. Features of existing manipulator technology.
- C. Scaling relationships for major arm components.
- D. Design of a particular small manipulator.
- E. Comments on future possibilities.

- 301 A Mechanical Arm Control System, Richard C. Waters, January 1974. (\$1.30) (AD-A004-872).

This paper describes a proposed mechanical arm control system, and some of the lines of thought which lead to the current design. It is divided into five main sections:

1. Some Ideas on Control
2. The Basic Capabilities of the Arm Controller
3. The Internal Structure of the Arm Controller
4. The Dynamic Level of the Arm Control System
5. The Procedural Level of the Arm Control System

- 308 Force Feedback in Precise Assembly Tasks, Hirochika Inoue, August 1974. (AD-A011-369). (\$1.30)

This paper describes the execution of precise assembly tasks by a robot. The level of performance of the experimental system allows such basic actions as putting a peg into a hole, screwing a nut on a bolt and picking up a thin piece from a flat table. The tolerance achieved in the experiments was 0.001 inch. The experiments proved that force feedback enabled the reliable assembly of a bearing complex consisting of eight parts with close tolerances. A movie of the demonstration is available.

- 319 Localization of Failures in Radio Circuits - A Study in Causal and Teleological Reasoning, Allen Brown, Gerald Sussman, December 1974. (AD-A011-839).(\$1.30)

This paper examines some methodologies for diagnosing correctly designed radio circuits which are failing to perform in the intended way because of some faulty component. Particular emphasis is placed on the utility and necessity of good teleological descriptions in successfully executing the task of isolating failing components.

- 323 Orienting Silicon Integrated Circuit Chips for Lead Bonding, Berthold K. P. Horn, January 1975. (In Computer Graphics and Information Processing, September, 1975). (\$0.90).

Will computers that see and understand what they see revolutionize industry by automating the part orientation and part inspection processes? There are two obstacles: the expense of computing and our feeble understanding of images. We believe these obstacles are fast ending. To illustrate what can be done we describe a working program that visually determines the position and orientation of silicon chips used in integrated circuits.

- 328 Heuristic Techniques in Computer Aided Circuit Analysis, Gerald Jay Sussman and Richard Matthew Stallman, March 1975. (\$1.30). (ADA-021171).

We present EL, a new kind of circuit analysis program. Whereas other circuit analysis systems rely on classical, formal analysis techniques, EL employs heuristic "inspection" methods to solve rather complex DC bias circuits. These techniques also give EL the ability to explain any result in terms of its own qualitative reasoning processes. EL's reasoning is based on the concept of a "local one-step deduction" augmented by various "teleological" principles and by the concept of a "macro-element". We present several annotated examples of EL in operation and an explanation of

how it works. We also show how EL can be extended in several directions, including sinusoidal steady state analysis. Finally, we touch on possible implications for engineering education. We feel that EL is significant not only as a novel approach to circuit analysis but also as an application of Artificial Intelligence techniques to a new and interesting domain.

- 329 Parsing Intensity Profiles, Tomas Lozano-Perez, May 1975. (\$1.30). (ADA-021172).

Much low-level vision work in AI deals with one-dimensional intensity profiles. This paper describes PROPAR, a system that allows a convenient and uniform mechanism for recognizing such profiles. PROPAR is a modified Augmented Transition Networks parser. The grammar used by the parser serves to describe and label the set of acceptable profiles. The input to the parser are descriptions of segments of a piecewise linear approximation to an intensity profile. A sample grammar is presented and the results discussed.

- 335 Image Intensity Understanding, Berthold K. P. Horn, August 1975. (\$2.10). (ADA-021135).

Image intensities have been processed traditionally without much regard to how they arise. Typically they are used only to segment an image into regions or to find edge-fragments. Image intensities do carry a great deal of useful information about three-dimensional aspects of objects and some initial attempts are made here to exploit this. An understanding of how images are formed and what determines the amount of light reflected from a point on an object to the viewer is vital to such a development. The gradient-space, popularized by Huffman and Mackworth is a helpful tool in this regard.

- 351 A State Space Model for Sensorimotor Control and Learning, Marc Raibert, January 1976. (\$1.30).

This is the first of a two part presentation of a model which deals with problems of motor control, motor learning, adaptation, and sensorimotor integration. In this section the problems are outlined and a solution is given which makes use of a state space memory and a piece-wise linearization of the equations of motion. A forthcoming companion article will present the results of tests performed on an implementation of the model.

- AI-TR-346 (Thesis) John M. Hollerbach, Hierarchical Shape Description of Objects by Selection and Modification of Prototypes, November 1975. (\$3.35).

An approach towards shape description, based on prototype modification and generalized cylinders, has been developed and applied to the object domains pottery and polyhedra: 1. A program describes and identifies pottery from vase outlines entered as lists of points. The descriptions have been modeled after descriptions by archeologists, with the result that identifications made by the program are remarkably consistent with those of the archeologists. It has been possible to quantify their shape descriptors, which are everyday terms in our language applied to many sorts of objects besides pottery, so that the resulting descriptions seem very natural. 2. New parsing strategies for polyhedra overcome some limitations of previous

work. A special feature is that the processes of parsing and identification are carried out simultaneously. With this descriptive approach, the evidently unrelated domains of pottery and polyhedra are treated similarly. Objects are segmented into multiple generalized cylinders. The cylinders are then described by assigning a prototype, a standard shape from a small repertoire, which is modified to conform more exactly with the cylinder. The modifications are structured hierarchically and specify the degree of modification as coarsely or precisely as desired. Some modifications are specific to a given prototype, others are applicable to several of them. The emphasis throughout this work has been to develop useful, qualitative descriptions which bring out the significant features and subordinate lesser ones. To this purpose curved lines representing the boundary of vases have been quantized into a few curvature levels. Line, region, and volume shapes are all described by assigning and modifying prototypes. In each instance the prototypes are specialized to the domain, and pose different problems in selection and modification.

Working papers relevant to productivity technology:

42. Mechanical Arm Control, Richard C. Waters, March 1973 (see A.I. Memo 301)
51. Finding Components on a Circuit Board, Tomas Lozano-Perez, September 1973
52. Tracking Wires on Printed Circuit Boards, Tim Finin, October 1973
54. A Scenario of Planning and Debugging in Electronics Circuit Design, Gerald J. Sussman, December 1973
59. GT40 Utility Programs and the LISP Display Slave, Michael Beeler, Joseph Cohen & Jon L. White, January 1974
61. Qualitative Knowledge, Casual Reasoning, and the Localization of Failures, Allen L. Brown, March 1974
64. Mini-Robot Group User's Guide, Meyer Billmers, March 1974
68. X-Y Table User's Manual, Noble Larson, May 1974
69. Kinematics of the MIT-VICARM Electronic Manipulator, H. Inoue and B.K.P. Horn, May 1974
71. Advice on the Fast-paced World of Electronics, Drew McDermott, May 1974
77. MAPPER Information, David Taenzer, September 1974.
91. Arrangement of Motions in a Robot Assembly Machine, Carl Flatau, November 1974.
96. Progress Report on Visual Inspection of Solder Joints, Dave Taenzer, April 1975.
98. Notes Relating to the Design of a High Quality Image Sensor, Berthold K. P. Horn, June 1975.
99. Kinematics, Statics, and Dynamics of Two-D Manipulators, Berthold K. P. Horn, June 1975.
100. The Application of Linear Systems Analysis to Image Processing. Some Notes., Berthold K. P. Horn, 1974.
106. Visual Tracking of Real World Objects, Glen Speckert, July 1975.
107. Transparency, Dan Stefanescu, July 1975.
116. The Fundamental Eel Equations, Berthold K. P. Horn, December 1975.
117. Computer Detection of Bent Fingers in Lead Bonding Frames, Walter L. Mitnick, January 1976.
118. Knowledge Driven Recognition of the Human Body, Glen Speckert, January 1976.