



STANFORD RESEARCH INSTITUTE
Menlo Park, California 94025 · U.S.A.

PLAN FOR A COMPUTER-BASED CONSULTANT SYSTEM

by

Nils J. Nilsson

J. Agin

B. G. Deutsch

R. Fikes

E. D. Sacerdoti

J. M. Tenenbaum

Artificial Intelligence Center

Technical Note 94

May 1974

The work reported herein was sponsored by the Advanced Research Projects Agency of the Department of Defense under Contract DAHCO4-72-C-008 with the U.S. Army Research Office.

CONTENTS

ABSTRACT	iii
LIST OF ILLUSTRATIONS	v
LIST OF TABLES	v
I INTRODUCTION	1
A. Project Goal	1
B. Applications	2
C. Demonstration System	3
D. Overview of Report	4
II DESCRIPTION OF DEMONSTRATION SYSTEM	5
A. Domain Selection	5
1. Scientific Value	6
2. DoD Relevance	6
3. Amenability to Sound Management	6
4. Extendibility and Generality	7
5. Experimental Convenience	7
6. Compatibility with Speech-Understanding Research	7
B. Description of the Domain	8
C. Apprentice Tasks	11
D. Capabilities of the Demonstration System	11
1. Dialog	11
2. Specific Capabilities	22
III RESEARCH PROBLEMS	29
A. Planning and Execution Monitoring	29
B. Vision	31
C. Natural Language	34
1. Discourse Semantics	35
2. Pragmatics and Model Interaction	36
3. Description, Generation, and Comprehension	37

III	RESEARCH PROBLEMS (Continued)	
	D. System Integration	41
IV	PROJECT PLAN	43
	A. Phasing	43
	B. The Phase I System	44
	C. Organization	44
	D. Research Tasks	45
	1. Planning and Execution Monitoring (Leader: Richard Fikes)	45
	2. Vision Research (Leader: J. M. Tenenbaum)	48
	3. Natural-Language Investigation (Leader: Barbara Deutsch)	52
	APPENDICES	
	A DESCRIPTION OF A SIMPLE AIR COMPRESSOR	55
	B ANALYSIS OF AIR COMPRESSOR CHECKOUT TASK	63
	REFERENCES	71

ABSTRACT

This report describes the goals and plans for a five-year project to develop a computer-based system that will serve as an expert consultant to a human apprentice. Together, the system and the apprentice will be engaged in a task of "checking out" and repairing electro-mechanical equipment in a workstation domain.

We have specified the capabilities of a system that we think is achievable in five years, and we have identified the major research problem areas involved. These include modeling and representation, automatic planning and problem solving, machine vision, natural-language communication, and system integration.

The demonstration system will give the human apprentice advice about how to diagnose equipment faults, how to repair them, and how to assemble and disassemble equipment. It will also give instructions about workstation procedures and information about such matters as the names of components and uses of tools. The level of detail of this advice and information will be automatically tailored to the individual needs of the apprentice, who will be able either to ask specific questions or seek general guidance. An important means of communication will be through natural language. The system will have the ability to understand continuous speech about the domain of interest.

Another important perceptual channel will be computer-based vision. The visual subsystem will be able to identify equipment components, inspect assemblies, and generally monitor progress at the workstation.

The demonstration system itself will serve to illustrate the feasibility of applying the constituent computer technology to problems of

equipment operation, maintenance, and repair; to remote site and vehicle support; and to similar applications.

We are proposing to divide the project into two consecutive phases. Phase I, to last three years, will be concluded by demonstrating a scaled-down system. This system will contain many of the planning and visual abilities of the proposed final system, but it will not have a continuous speech-understanding component; it will operate at a rate slower than real time, and its domain of expertise will be limited to a single version of a simple air compressor. During this first phase, we will be conducting the research necessary for constructing such systems, testing out various approaches, and further elaborating our plans for the final demonstration system. During Phase II, the last two years, we will use the results of the first demonstration to adjust our research and design strategies. For the final demonstration system, we will relax the restrictions of Phase I. The system will have become expert about a class of equipment including pumps, motors, and air compressors, and it will accept spoken natural language as input. Our report gives full details on the phasing of the project and specific plans for the research to be conducted.

ILLUSTRATIONS

1	Proposed Workstation	10
2	An Air Compressor	12
A-1	Parts List Automatic Switch Control, Storage Tank Type	60
A-2	Parts List Pump Unit	61

TABLES

1	Workstation Components	9
---	----------------------------------	---

I INTRODUCTION

A. Project Goal

This report describes a research program the goal of which is to develop a computer system that can serve as an expert consultant to a human about a domain of knowledge. As just one example, such a system could advise and guide a human apprentice about the operation, assembly, fault diagnosis, and repair of electromechanical equipment.

Consider the following excerpt from a dialog between an expert and an apprentice during checkout of a small air compressor:

Expert: Have you checked the belt?

Apprentice: What belt?

E: I am pointing at the belt housing frame. The belt is inside. To check it, you will have to remove the belt housing cover.

A: Where's that?

E: It's attached to the back of the belt housing frame by ten screws.

A: O.K. I have the cover off. The belt is hanging loosely around the two pulleys.

E: Please show me the belt. Thank you. Yes, the belt needs to be replaced.

A: I have a new one here. I'll replace it.

E: O.K., but first you should check some other things.

The goal of this project, simply stated, is to produce a computer system that can fill the role of an expert such as the one in the above dialog both in knowledge of the domain and in ability to communicate that knowledge. Later in this report, we shall analyze a more extensive dialog to extract the research problems confronting the development of a

computer-based expert. We can state briefly here that the major problems are involved with modeling and representational techniques, natural-language communication, computer perception of visual images, automatic planning and problem solving, and system integration.

B. Applications

Why is it important to develop computer-based consultant systems? First of all, even a trained person has frequent need of consulting manuals, references, and other human experts, so there is a definite need for readily available expert knowledge. In both industry and the military, equipment maintenance and training of maintenance people represents a large budget item. Any technology that can reduce these expenditures and lessen the need for utilizing scarce human experts will obviously be greatly in demand. We are also convinced that the need for consultation cannot be satisfied merely by writing more and better manuals. A sophisticated computer system seems to us essential. We might cite some advantages that such a computer system would have over written (or prerecorded voice) manuals. It would be:

- More closely adapted to the particular user and his needs (e.g., the level of detail can be changed dynamically to match the expertise of the user).
- Richer in human interface possibilities (e.g., speech communication in both directions).
- More easily modifiable.
- Able to reproduce up-to-date copies of the current versions easily, and to distribute them via computer nets.
- Able to monitor the progress of the user at his task.
- Able, as an active partner, to take the initiative when it is appropriate.

Computer-based consultant systems would be useful in several application areas. The following come easily to mind.

- Systems that are expert in equipment design, operation, maintenance, assembly and/or repair.
- Systems that are expert in the operation and procedures of specialized remote sites and vehicles, such as
 - Radar stations
 - Submarines and other naval ships
 - Manned satellites
 - Aircraft
 - Polar stations.
- Medical systems:
 - Diagnosing
 - Surgical consulting.
- Air traffic control systems.
- Management support systems.*
- Military command and control systems.
- Legal reference systems.
- Office information systems.
- Computer programmer's adviser.

In these applications, the emphasis would be on building systems that can converse naturally with humans and aid them by providing advice, answering queries, and so on.

C. Demonstration System

Before large-scale computer-based consultant systems for these applications will be feasible, there are a number of research problems that must be solved. We are proposing here a highly focused five-year research

* An ongoing SRI project (Contract N00014-71-C-0210) sponsored by the Office of Naval Research is already beginning to make use of artificial intelligence (AI) techniques in research on large-scale computer-based management systems for naval repair operations.¹ (References are listed at the end of this report).

project designed to attack these problems. At the end of this period, specifically 1 April 1978, we plan to demonstrate a model computer consultant system. This demonstration system will show the feasibility of our solutions to the major research problems and will serve as a prototype from which actual applications systems can be developed.

D. Overview of Report

In the next section we will describe the domain of expertise selected for the five-year demonstration system and the criteria used in the selection of that domain. Next, we shall describe the planned capabilities of the demonstration system and discuss the major research problems that arise in attempts to meet them. Finally, we shall present our plan for attacking these research problems and for integrating the results into the demonstration system itself.

II DESCRIPTION OF DEMONSTRATION SYSTEM

A. Domain Selection

The domain we have selected for the demonstration system is a workstation in which various types of electromechanical devices are to be checked out and repaired by a human apprentice. The system will be able to answer queries from the apprentice about the properties and purposes of tools, subassemblies, and parts. It will be able to provide advice and instructions on how to assemble, disassemble, operate, troubleshoot, and repair the devices being maintained at the workstation. Although we expect the demonstration system to be competent with a range of electromechanical devices, we shall concentrate attention at the start specifically on an air compressor. Such a choice allows us, even at the beginning, to deal with a reasonably complex system consisting of an electric motor, a pump, air storage tank and lines, and a feedback control device.

In this section we will describe the criteria we used in selecting repair of electromechanical devices as the domain for our project, and then we will give a more detailed description of the workstation and the capabilities that we expect our demonstration system to have.

We used six general criteria in evaluating possible domains.

- Scientific value
- DoD relevance
- Amenability to sound management
- Extendability and generality
- Experimental convenience
- Computability with speech-understanding research.

1. Scientific Value

Will the demonstration system require and will it stimulate solutions to the major research problems confronting the development of computer-based consultant systems? The domain chosen must be an appropriate area in which to conduct the research needed to solve these problems.

The domain must be complex enough to render infeasible the use of ad hoc methods that would not be applicable to large domains (such as storing canned answers to all the expected questions). It seems clear to us that most application domains are complex enough to require that knowledge about them be stored in general models and procedures from which answers and advice must be computed. But the domain must also be well bounded; we do not want it to merge so continuously into ever larger and more complex problems that we would have difficulty at any stage in limiting its scope. In summary, we want to be working on a project that has a good chance of making fundamental scientific contributions.

2. DoD Relevance

Will the demonstration system have obvious relevance to the Defense Department? Although we don't insist that the system itself be operationally useful to DoD problems, we want it to be clear that the technology of the demonstration system could be applied in actual situations with only modest additional development (not research). Related to this condition is the requirement that the system not be "toy-like."

3. Amenability to Sound Management

Is the demonstration system compatible with developing a good management plan? First, can we define our objectives precisely enough so that we can tell at the end of the project how well we achieved them?

Second, can we develop clear subgoals and milestones that will help us achieve the end goal?

4. Extendability and Generality

In addition to being relevant to specific DoD problems, we require that our research lead to results that are broadly applicable. We want to develop a system that can be extended or generalized to a wide family of applications. Thus, we think it would be a mistake to choose a domain in which highly specialized or esoteric knowledge is required. Rather, we favor a domain requiring large amounts of knowledge of the "common sense" variety. With such a choice, we think we can maximize the likelihood of developing fundamental technology.

5. Experimental Convenience

Is the demonstration system convenient experimentally? Are any highly specialized facilities (other than the computer and associated input/output equipment) required? Is any specialized knowledge required that might be difficult for us to obtain? Is the system on a scale suitable, both physically and conceptually, for the laboratory? (On these grounds, for example, we would rule out research experiments with aircraft engines--too large--and wristwatches--too small.)

6. Compatibility with Speech-Understanding Research

We are convinced that natural language input/output is crucial both to the demonstration and to research progress. Thus, we would like to choose a domain whose vocabulary and semantics are compatible with current ARPA-supported projects in speech understanding, especially the project here at SRI.

B. Description of the Domain

The workstation will contain, besides the devices being worked on, a set of necessary tools such as wrenches and screwdrivers, a workbench with vise, a cabinet, drawers and pegboard for tool storage, electrical outlets, miscellaneous spare parts, lubricants, and such other items as we think useful. A list of typical workstation components is shown in Table 1. This list may be expanded and modified as the research progresses, but the number and complexity of items to be used in the demonstration system will be no less than is indicated in the table.

Although the physical layout of the workstation has not yet been completely designed, the general idea of what we have in mind is shown in Figure 1.

We plan to mount in the workstation a television camera system that will be able to see the equipment being worked on, the workbench, and the tool storage area. The camera will be augmented by a low-power laser range finder. This visual subsystem will be used by the system to monitor the progress of operations in the workstation. It will also serve as an input channel that can be used by the apprentice for queries to the system. For example, the apprentice may ask the system to name the part he is holding in front of the camera. Or, for example, in response to a system suggestion to the apprentice to remove the pump, the apprentice may point to a part and ask the system, "Is this the pump?"

The laser beam projected by the range finder will also be utilized as a pointer to designate things in the work station. With this beam, the system will be able to point out various parts, tools, and other objects to the apprentice.

A display console will also be installed in the work station. The system will be able to display sketches of parts and subassemblies to the apprentice. Even though such a display system could be quite useful,

Table 1

WORKSTATION COMPONENTS

The workstation room and utilities (lighting, electric outlets, clock)

Workbench and stool, vise

Storage areas for tools (drawers, cabinets, boxes, and pegboard)

Spare parts needed in repair

Storage areas for spare parts

Tools and equipment needed for checkout and servicing*

- Socket wrench set consisting of
 - 6- and 12-point sockets and drivers
- Open end wrenches
- Box end wrenches
- Combination wrenches
- Allen wrenches
- Phillips head screwdrivers
- Ordinary screwdrivers
- Wheel puller
- Punch and chisel set
- Nut drivers
- Pliers

Lubricating oil

Solvent or cleaner

Rags

*Partial list only.

we think that we will not fully exploit the use of graphics, at present, since we do not believe that graphics should be a primary concern of the project. Instead, we plan to make extensive use of the laser pointer to reduce the need to display objects pictorially. The primary means of communication between the knowledge system and the apprentice will be natural language--specifically speech. Thus, a microphone and loud-speaker will constitute a principal input/output channel.

PROPOSED WORKSTATION



SA-1530-41

FIGURE 1 PROPOSED WORKSTATION

C. Apprentice Tasks

Work on the project has begun by considering a specific device to be maintained in the workstation--namely, a small air compressor (shown in Figure 2 and described in detail in Appendix A).

A typical main task for the apprentice in the workstation is to check out the air compressor and prepare it for operation. This task will involve several subtasks, including


- Operation of the compressor and its components
- Inspection and test
- Diagnosis of faults
- Replacement and/or repairs of components
- Assembly and disassembly of the compressor and its components
- Service operations such as lubrication and cleaning.

In Appendix B we present in some detail an analysis of the air-compressor checkout task. This description elaborates some of the subtasks that will be faced by the apprentice and indicates how the demonstration system can aid him.

D. Capabilities of the Demonstration System

1. Dialog

Before listing the specific capabilities of the system, we want to present transcribed excerpts from a tape-recorded dialog between a human expert and apprentice. This dialog serves as the basis for indicating what kinds of knowledge and abilities the system must have. It also defines our ultimate goal--to be able to duplicate the performance of the human expert. We do not expect our five-year demonstration system to attain that goal, but we do expect it to be capable of participating in a dialog much like the one presented here. In the sections following



SA-1530-70

FIGURE 2 AN AIR COMPRESSOR

the dialog, we describe explicitly the anticipated capabilities of the five-year system.

We have provided comments with the dialog to elaborate on what the expert is doing at each stage. These comments are intended to provide insight into what abilities will be needed by a computer-based consultant.

During the experimental session that produced this dialogue, the expert was shielded visually from the apprentice and the compressor, so that all communication was through speech. This experimental situation served to elicit interesting verbal descriptions from the expert and also to illustrate how the lack of visual communication severely restricts the expert's ability to guide the apprentice and monitor the performance of the task.

We present a rather long section of the dialog to illustrate the many capabilities that must be possessed by such a system. The reader may want to scan over this section quickly on a first reading.

The dialog begins with the compressor disassembled:

<u>Dialog</u>	<u>Comments</u>
Expert: The task is to re-assemble the compressor.	The expert provides an overview.
Apprentice: Uh huh.	
E: The first thing we'll do if to attach the motor to the platform.	The expert determines a sequence of operations (i.e., a plan) that pertains to the task.
A: Attach the motor to the platform? O.K. And now, let's see. This looks like..., this must be a motor, and um.... Let's see, that's number 17 there, and number 17 is motor or gas engine. O.K.	The apprentice is looking at a diagram with parts numbered. (The laser pointer would have helped here; the expert could merely have pointed at the motor and platform. Alternatively, vision would have made it possible for the apprentice to point and ask for verification.)

<u>Dialog</u>	<u>Comments</u>
E: Yes, and it should be placed on the platform, and it is connected to the platform by four bolts.	Expert realizes that apprentice has identified and found the motor.
A: Now, where on the platform should it go?	
E: Near the left side. You can tell the position from its previous position by where there's no paint.	The expert answers a question that provides more detail about the operation. Location specification is such that a human can easily understand it. (Laser pointer would have helped.)
A: Ah hah!	
E: The shaft of the motor should point toward the wall and should extend through the belt housing frame.	The expert provides still more detail in response to the apprentice's question. He makes the position unique. (Vision would have allowed the expert to verify correctness.)
A: Right. O.K. Now it's... I have it positioned correctly.	
E: O.K. Now find a set of four bolts. Each bolt has a nut and three washers attached to it.	The expert realizes that the motor location difficulty is resolved. Expert gives a "distinguishing-features" description of the bolts so that the apprentice can identify them among the other parts on the workbench.
A: O.K.	
E: Be careful about the order of washers. There should be one washer between the bolthead and the motor.	The expert is aware of common errors that humans make and, therefore, gives a warning.
E: Yes.	The expert must be able to understand and respond to the apprentice's descriptions of where the washers should go.

<u>Dialog</u>	<u>Comments</u>
A: One is underneath the uh...the uh...little platform that the motor sits on but on top of the.... They're all three in different places.	
E: No, that's not correct.	
A: They're all together?	Meaning they all go together.
E: The first one is between the head of the bolt and the motor.	The expert begins a more detailed description of the bolt assembly to alleviate the difficulty.
A: Yes.	
E: And the other two are below the platform.	
A: Oh, the other two are completely below the platform.	
E: Yes, the smallest one, which is a lock washer, is next to the nut.	
A: I see.	
E: The purpose for those washers is that the motor has to be able to slide back and forth to adjust the tightness of the belt. The two large washers are to facilitate that, and the small washer is a lock washer which keeps the nut from coming loose.	The expert gives a functional description of the washers to increase the apprentice's understanding of the assembly. Such an understanding should allow the apprentice to further verify that the assembly has been made correctly.
E: Are you working on the second bolt?	The expert must maintain a model of the current situation so that he can respond to difficulties and to questions. (Vision would help here, although this would be a hard vision problem.)
A: Yes, I'm on the second bolt now.	(Vision could be used here to check.)
E: O.K.	

<u>Dialog</u>	<u>Comments</u>
A: O.K. All four bolts are now in place and loosely fastened.	
E: O.K. Would that task have been significantly easier if the belt housing frame had been removed?	Since this was a difficult step, the expert asks for feedback from the apprentice. In this case, the expert makes a hypothesis for the apprentice to comment on.
A: Which is the belt housing frame?	
E: It is the vertical piece of sheet metal at the rear of the platform.	This is a descriptive and positional definition. The system must be able to describe parts for identification. (Laser pointer would have helped.)
A: Well, not significantly, no. I don't think it would have made too much difference. It would have been easier if the platform that the motor sits on were removed from the cylinder at the bottom.	"The cylinder at the bottom" is a description of the tank.
E: Yes, but the platform and the cylinder are welded together and, therefore, cannot be disconnected. Can you give me any advice that I could give in the future to a person doing that task?	The expert understands the apprentice's comment, recognizes the impossibility, and asks for any advice.
A: Oh yes, I found on the final bolt that it is easier to keep both hands underneath the platform and let the weight of the bolt push it through the washers. So, it was better to leave both hands underneath the platform so that I had two hands to manipulate the washers and bolt with.	To understand this suggestion, the expert would have to know that the phrase "on the final bolt" tells when the apprentice found an easier way of doing this task. It does not indicate that the suggestion is relevant only to the final bolt.

<u>Dialog</u>	<u>Comments</u>
E: O.K. Thank you. Is there any special ordering of the bolts that would be significantly better than any other?	The expert realizes he has given the apprentice an unordered set of operations; he asks if order is important.
A: No, I don't think it would have made any difference.	"It" refers to "a special ordering of the bolts."
E: O.K. Thank you. Next we want to attach the wires to the motor. There is a cable coming out of the pressure switch.	Each assembly operation begins with the expert helping the apprentice to locate the parts to be connected. This is followed by a series of steps to achieve that connection.
A: Yes.	
E: And it has two orange-colored wires at the end.	The expert uses color as the distinguishing feature. (The laser pointer would have helped.)
A: Yes.	
E: Those are to be inserted through the hole in the side of the rear of the motor.	The expert uses the relative terms "side" and "rear."
A: Uh huh.... I see it.	
E: Each wire is to be attached to a bolt.	
A: Bolt? Yes. Does it matter in which order we uh... which one goes to which?	
E: No, it doesn't.	
A: Ah! Good. Now, should I unscrew the nuts from the bolts?	
E: No, the wire goes on top of the nut.	The expert should be able to answer questions where the apprentice suggests an operation to be done and the expert approves or disapproves.

<u>Dialog</u>	<u>Comments</u>
A: I see.	The expert erroneously assumes here that the apprentice has put the wires in place over the bolts. (Vision would have helped avoid this error.)
E: There is a pair of other nuts.	
A: Other nuts!	
E: There is a pair of nuts that are sitting on the desk. They are close to a small plate about three inches long.	The expert must help the apprentice find all the parts needed to complete the assembly operation. Location information is given by the expert, indicating what is close that may be easier to find. (The laser pointer would have helped.)
A: Wait a minute--should I be looking for the nuts, or should I.... I haven't put the thing in place, yet.	
E: O.K. Put the wires in place so that the loops are over the bolts.	The expert corrects his model of the current situation.
A: O.K. They're both in place.	
E: O.K. Now we want to find the two small nuts.	
A: O.K. I've found two nuts, but they seem to have screws next to them.	
E: Yes, that's O.K. Those screws go into something else.	The expert recognizes that since screws and nuts are together, the apprentice assumes them to be a set. The expert makes a comment to correct that assumption. (Vision would have helped.)
A: Those are different screws.	
E: Yes.	
A: All right, I've found two nuts.	

<u>Dialog</u>	<u>Comments</u>
E: I suggest you start them by hand.	The expert could have assumed that the apprentice could attach the wires without further assistance, but by now he has a rough model of the apprentice's level of expertise and decides that detailed instruction is needed.
A: All right. One of them's on,... and now number two is on.	(Vision could be used to verify this.)
E: Now, you need to tighten them using a nut driver.	
A: A nut driver! What's a nut driver?	
E: A nut driver looks like a screwdriver and is in the yellow plastic case leaning against the wall.	The expert knows what tools should be used for each operation and can help the apprentice identify and locate tools. Descriptions are by analogy, color, and position. (Laser pointer would have helped.)
A: The yellow plastic case... Oh, I see the yellow plastic case. Yes.	The expert knows that handle color uniquely identifies which nut driver.
E: The driver is 11/32 inch.	The expert uses very detailed specification to verify acquisition of the correct tool.
A: Yes, it says 11/32 right here on the handle, Number 11.	
E: O.K., you can use that tool to tighten the nuts.	
A: O.K.	
E: Now, the loop at the end of each wire should be over the bolt.	Since this operation has involved several substeps, the expert initiates a checkout sequence before proceeding. (Vision might have helped here.)

<u>Dialog</u>	<u>Comments</u>
A: Uh-huh.	
E: And it should be between the nut that you are tightening and the nut that was already on the bolt.	
A: I understand. Should it be very tightly tightened? Or just, uh, snugly?	"Tight" and "snug" are fuzzy terms. Note that vision could <u>not</u> help here.
E: No, just snug.	
A: All right. We're done.	
E: O.K.	
A: Should I put the nut driver back?	
E: Yes, please. Now we want to attach the plate to the back of the motor. The plate covers the wire connection that you just made.	The expert needs to keep track of the location of each tool in the workshop. Expert relates next operation to the previous one.
A: Uh-huh.	
E: The plate is approximately 3 inches long and one inch wide.	Again, the expert helps the apprentice find the parts involved in the assembly step. (Laser pointer would have helped.)
A: Oh yes! I see it.	
E: O.K. and you use...	
A: It says lubrication on it-- do not overoil?	
E: Yes, that is it.	Expert is interrupted by a question from the apprentice, recognizes it as a verification of the plate identification, and responds accordingly. (Vision could have enabled verification without such a lengthy dialog; the apprentice could have held the piece up for the expert to see.)
A: Ah-hah!	

<u>Dialog</u>	<u>Comments</u>
E: And you should use the two screws that you mentioned earlier.	Expert refers to a past event to help the person locate the screws.
A: I see.	
E: Now, slide the pulley onto the shaft and tighten the screw down.	
A: I should hand-tighten it.	
E: No, use an Allen wrench.	
A: An Allen wrench? What's an Allen wrench?	
E: We have a combination Allen wrench. First of all, an Allen wrench is a hexagon-shaped rod that will fit on the head of that screw, and the particular tool that we have looks like a pocket knife, since it has...	The expert uses an analogy to describe the tool and recognizes the interruption from the person finding it. (Laser pointer would have helped locate the set of wrenches.)
A: Oh. I see...just like a pocket knife...now....	
E: O.K. Is it true that the pulley is on the shaft?	Again, the expert asks questions to update its model and to cause the apprentice to go through a checkout procedure. (Vision could have verified this.)
A: Yes, the pulley is on the shaft.	
E: Yes, and the screw is tightened?	
A: The screw is tightened.	
E: The screw is next to the flat part of the shaft?	
A: Um, yes, the screw is on the flat part of the shaft.	
E: O.K. Now, we need to work on the pump.	

<u>Dialog</u>	<u>Comments</u>
<p>A: There's a problem coming up. There's a...on the... uh...this part that's sticking up in back that has a protrusion. Oh, perhaps my shaft is.... Oh, perhaps there's another shaft. There's a fat shaft and a thin shaft. Oh yes, those seem to fit better with the... uh.... Well, let's see, should the fat shaft point toward the wall or should the thin shaft point toward the wall?</p>	
<p>E: I'm not sure I understand the description. The shaft that should point toward the wall is the longest protrusion from the pump, and it has a small half-moon shaped piece of metal sitting near the slot.</p>	<p>The expert need not understand everything the apprentice says. In this case, the expert understands the nature of the problem and can therefore respond appropriately.</p> <p>The apprentice and the expert are using different descriptors for the shaft; e.g., "fat" and "thin" versus "long protrusion" and "has a small half-moon shaped...."</p>
<p>A: So, the longest protrusion should go toward the wall.</p>	<p>In this case, the apprentice alters his descriptors. (With vision the expert would have understood fat and thin and this problem would have been avoided.)</p>
<p>E: Yes, that's correct.</p>	

2. Specific Capabilities

In analyzing the preceding dialog, we can group the capabilities of the expert under the following headings: planning and execution monitoring, vision, and natural language. In the remainder of this section, we shall list the specific capabilities of the five-year demonstration

system as regards these topics. Then, in the next section, we shall discuss the research problems connected with them.

a. Planning and Execution Monitoring

The expert system will have the following capabilities for planning and execution monitoring:

- (1) The basic capability to generate step-by-step plans (i.e., sets of instructions) for converting a device (e.g., the compressor) from one arbitrary state of assembly into another state. This capability will exist at several levels from major components (e.g., removing the pump from the compressor) to the simple fasteners (e.g., removing the connecting pin from the pump piston).
- (2) The capability to generate plans in a hierarchical fashion in order to give them to the apprentice at a level of detail that matches his expertise. The system will, upon request by the apprentice, provide more detailed instructions for any step in the reconfiguration process.
- (3) The ability to answer specific questions about instructions it has given, questions like, "What is the purpose of this step?" or "What kind of wrench should be used for this step?"
- (4) The capability to give technique-related advice when the apprentice is having difficulty with a step. Such advice might be, "Tighten the nut with your fingers before using the wrench" or "Pulley alignment can be tested by placing your hand across the front of both pulleys."

- (5) The capability to give the apprentice step-by-step instructions for diagnosing a malfunctioning device (i.e., trouble-shooting). The system will be sufficiently knowledgeable about the nominal behavior of the device to be able to recognize and analyze abnormal behavior. It will possess general knowledge about the kinds of malfunctions. The apprentice will be asked to perform tests on the device to assist in the generation and verification of the malfunction analysis.
- (6) The capability, once the trouble-shooting is successfully completed, to guide the apprentice through a repair procedure. Often, the repair procedure will be simply to replace a component, in which case it will be equivalent to a disassembly-assembly reconfiguration procedure.
- (7) The capability to describe the use of the tools found at the workstation. Hence, the system will be able to tell the apprentice which tool to use for a step, help to identify and locate the tool, and give instructions on its operation.
- (8) The capability to give warnings and reminders to help the apprentice avoid common errors and to alert him to irreversible or dangerous steps. Similarly, the system will be able to inform the apprentice about such things as the criticality of a step, the accuracy to which an adjustment must be made, and the tightness required for a fastening.
- (9) The capability to monitor and inspect the apprentice's work to ensure that the operation is proceeding nominally. When the system becomes aware of an error or of an unexpected event (from

interactions with the apprentice or from the system's sensing devices), it will alter instructions to the apprentice to deal effectively with the new situation.

- (10) The capability to take advice from the apprentice (or other humans). This task, in keeping with our basic premise that the system be a repository of expertise, appears to be a difficult one, but we can envision the acceptance of advice in at least two modes. First, the apprentice could answer specific questions from the system. (For example, "Can you suggest a preferred order for removing the motor-mounting bolts?" or "Would it have been easier to install the motor-mounting bolts if the belt housing had been removed?") Second, the system could elicit advice from the apprentice and simply store it in the form received. This advice could then be repeated to another apprentice having difficulty with the step, or it could be examined at a later time by the system designers for possible incorporation into the system's knowledge structure.

b. Vision

The vision subsystem will be capable of rendering a range of services of at least the scope and difficulty of the ones listed, along with samples of specific questions, below:

- | | |
|-------------------------|---|
| (1) Part description | "What color is the pump?"
"What are the dimensions of the washer?" |
| (2) Part identification | "What part or parts are lying on the table?"
"Is this a box wrench?"
"What is the laser pointing at?" |

- (3) Visual inspection "Is the hole too large?"
 "Is the cord frayed?"
 "Is this the correct part?"
 "Is the belt tight enough?"
 "Has the assembly been
 accomplished correctly?"
- (4) Performance monitoring "Tell me when the
 wrench handle reaches
 a horizontal position."
 "What is the reading on
 the pressure gauge?"
 Has the last bolt been
 installed yet?
- (5) Part and tool location "Where is the hammer?"
- (6) Graphic input/output "Point the laser at
 the upper left corner
 of the belt housing
 frame."
 "Display an exploded
 view drawing of the
 motor."
 "Adjust the TV camera
 to obtain a closeup
 view of the flywheel
 and fan belt."
 "Outline the pressure
 switch on the display."
- (7) Assistance in planning "Is there a space
 on the table in which
 we can set down the pump?
 If not, what must be
 moved?"
 "Will the grommet fit
 into the hole?"
 "Is there enough space
 to remove bolt X without
 removing the flywheel?"

c. Natural Language

Our goal is a system that understands connected speech at least as well as will the ARPA 1976 Systems.² More specifically, we

expect to handle task-related utterances in that subset of ordinary English expected to occur in the workstation environment in a few times real time. The utterances will contain words selected from a 1000-word vocabulary, and they will consist of syntactic constructions for the relevant subset of English. The number of levels of embedding and the kinds of conjunction and ellipses allowed will be limited. These limits will not be chosen arbitrarily. We will study protocols to determine how effort should be spent to extend our capability to accommodate natural English and where limits are likely to be easy to impose.

We do not expect to be able to handle false starts, but we do hope to handle short intra-sentence pauses. The system will not be able to learn new words. Input will be from a good-quality microphone in the workstation environment. How close we come to this goal depends, of course, on the progress achieved by the ARPA speech-understanding projects, especially in work on the acoustic and phonological parts of their systems.

The speech output will be relatively unsophisticated linguistically, although we do expect the system will be able to generate proper kinds of simple reference. Our efforts in this area will be concentrated on procedures and representations necessary for generating descriptions of objects and locations that are easily understood by people.

III RESEARCH PROBLEMS

The capabilities that we have just discussed are still beyond the state of the art in computer science, but it seems to us that they can be achieved in five years by a well focused and well planned research project. Here we shall discuss some of the research problems confronting the design of the demonstration system. They fall under the same three headings of the last section--namely, planning and execution monitoring, vision, and natural language. To these we will add a fourth, dealing with the special problems of integrating these abilities into a smoothly running system.

A. Planning and Execution Monitoring

The major problems here concern the automatic generation and execution of procedures for troubleshooting or reconfiguring the equipment. These procedures, or plans, must be constructed at varying levels of detail to match the apprentice's capabilities, and the system must "understand" the steps of the plan so that it can answer questions, monitor progress, and provide new plan steps when unexpected situations develop.

A good deal of previous work has been done on this subject, and it can be used as a base. We might mention specifically the following relevant work in automatic planning:

- The STRIPS planning system and its extensions.³⁻⁵
- The QA4 robot planning programs.⁶⁻⁸
- The SHRDLU system for planning in the BLOCKS world.⁹
- The HACKER system for generating and debugging plans in the BLOCKS world.¹⁰

Many of the problems not yet adequately handled in previous work are discussed by Fikes, Hart, and Nilsson.¹¹ Especially important is the need to develop a hierarchically-based planning system that can operate at varying levels of detail. Nilsson¹² and Sacerdoti⁵ have designed simple hierarchical systems, and these designs will be used as prototypes for the development of a system that can plan effectively in the complex environment of the workstation.

Our system will need to develop plans that are more complex than simple sequences of steps. For example, we will need various kinds of loops and conditionals ("fasten bolt 1, bolt 2, bolt 3, and bolt 4 loosely," or "loosen the set screw until the pulley moves freely on the shaft"); sets of unordered steps ("tighten the four bolts"); and sets of steps with a preferred, but not required, ordering ("remove the ten screws; you may find it easier to remove the bottom row before the top row"). An important research task is to develop a representation for plans that will allow us to generate, execute, and monitor such complex operations.

The workstation environment will require our system to be designed so that generation and execution of plans will be interwoven in a much more dynamic fashion than has been the case with previous planning systems. For example, questions from the apprentice about a particular step during plan execution may require the system to generate a new detailed plan that elaborates the step in question. Unexpected events or apprentice mistakes during plan execution may require the generation of new plan steps to get the operation back "on the track." Also, plans may include information-gathering steps whose results affect the course of the operation to the extent that planning cannot proceed beyond them.

The system will also be able to answer questions about a plan at any level of detail. To do this we will need (again) efficient ways of representing plans that reveal the purposes, preconditions, and results of each of the steps in the plan.

Fundamental to all of the capabilities we have discussed so far is the need for a rich hierarchical model of the workstation domain including the equipment being worked on and the apprentice himself. A major research task is the development of representations for this model that will allow the system to store and effectively use the information that an expert consultant is expected to have.

Some of the procedural information in these models will essentially be "canned" plans for doing common detailed operations such as tightening a nut-and-bolt assembly. During plan generation these stored procedures will be instantiated and used as detailed expansions of higher level planning steps. Inefficiencies and inadequacies may appear in detailed plans produced by the concatenation of such stored procedures, and the plan generator must be able to optimize and debug these plans appropriately. Some new work by Sacerdoti on developing special programs to criticize plans may be useful for this problem, as may Sussman's work on HACKER.¹⁰

B. Vision

The vision subsystem may be viewed as a set of information-gathering routines that may be called by any part of the consultant system to obtain specific kinds of information about the workshop environment. This subsystem is a specialized consultant system in itself, including capabilities for question answering and problem solving. What makes it necessary to consider the vision subsystem as somewhat distinct from the general problem-solving effort are the specialized requirements for visual data acquisition and image processing, and the representation of semantic concepts such as shape, color, and visual appearance. However, many of the techniques we propose for perceptual strategies are applicable to the rest of the system.

The vision-related research issues we will attempt to solve are in two broad areas: utilization of three-dimensional visual data and generation of perceptual strategies by interactive or automatic means.

Our first research objective will be to investigate the resolution and accuracy obtained by our laser range finder. The feasibility of obtaining range by phase measurement on a modulated laser beam has already been demonstrated. However, the limiting accuracy of the system is unknown, as are the tradeoffs between modulation frequency (wavelength), linearity, resolution, noise, and limiting range. These questions should be resolved through continued hardware development of our system.

A critical research objective is the development of adequate representations of curved three-dimensional objects that are distinguished primarily by structural shape descriptions. We do not necessarily envisage a single comprehensive formalism for describing shape but will continue throughout the course of this project to develop representations for more and more complex structures. Among the desired characteristics of a set of representations are the following: They should be capable of representing an arbitrary degree of detail without making the retrieval of the coarser features unwieldy; they should be capable of relating to models of surface characteristics including reflectivity, color, and texture; they should be capable of characterizing relationships among parts in an assembly; and above all, they should enable the system to produce (understand) descriptions easily understood (produced) by humans.

The algorithms by which the appropriate representations are abstracted from visual and range data are an integral part of the research on representation. The set of available representations and the algorithms for obtaining and transforming these representations we refer to as our "library of visual primitives."

For the first few years, we will concentrate on representing a fixed set of specific parts. These representations will be "model-based" in the sense that they refer to the precise geometric details of specific prototype parts. As the research proceeds further, we hope to become more and more "semantic-based," so that the features by which we recognize, say, a flywheel, do not include the number of spokes, but rather characteristics associated with most flywheels (It is round and its mass is concentrated toward the rim). An eventual goal is to use general knowledge about, for instance, compressors in order to analyze a compressor the system has never encountered before and to correctly identify its parts based on their functions.

To date, most scene analysis programs have been written to accomplish a limited objective using a small number of specialized data extraction techniques applied according to a fixed set of rules. No such program is sufficiently general to accomplish the variety of objectives listed above. Rather, a general vision capability must be based on a strategy program that can select from a large repertoire of specialized techniques those best suited for answering a particular question.

The techniques available for answering a specific request may range from a simple table lookup to a major analysis of an entire scene. In addition, the system can request the apprentice to aid in visual acquisition and interpretation. The visual subsystem will have a planning capability that will make use of knowledge about the context of the request and about the capabilities and limitations of the techniques available in order to arrive at an efficient way of answering the request. Other subsystems can advise the vision subsystem on how best to obtain the information they require.

Until our library of visual primitives is sufficiently rich to describe compressor parts, we will continue current work on perceptual

strategies in the domain of room scenes. Research topics will include the interactive specification of recognition strategies, the automatic generation of recognition strategies from features which an operator considers to be important, the representation of knowledge about a system's own capabilities, and special techniques for rapid identification of objects in a known context.

In the early stages of the program, recognition strategies for compressor parts will be hand-coded into the system and based on "distinguishing features," a set of tests that may economically distinguish among parts in a limited context. Later, we hope to be able to interactively "teach" the system which features or characteristics of an object may be reliably used for identification.

Finally, we are studying the feasibility of using a constraint satisfaction approach to narrow the range of possible interpretations for parts of a scene. Many constraints govern relationships of parts to each other in an assembly. These constraints arise from consideration of the function or operation of the parts and from methods of fabrication or assembly. For example, a motor must have a power cable connected to it; the absence of any such part connected to a particular component makes it unlikely that the component is a motor. It is hoped that a sufficient number of constraints will be discovered so that elimination of impossible relationships leaves only a small number of consistent interpretations for a scene.

C. Natural Language

Our system will be one of two participants in goal-directed dialog. To be understood, a sentence occurring in such a dialog must be interpreted in the context in which it appears. Both the discourse context (the actual dialog occurring so far) and the situational context (the

current workstation environment) are important. Thus, the language system will have to incorporate information from the workstation, an understanding of the workstation tasks, and a memory of recent events. It will have to be able to obtain information via vision and planning. Hence, the major emphasis in natural language will be in three areas: discourse semantics, pragmatics and model interaction, and description generation and comprehension. We will be attempting to answer such questions as

- What information is available from the task model and the state of the dialog to guide understanding?
- How should this information be structured?
- How can it be incorporated into a system that understands natural language?
- How can descriptions generated by people in task-oriented situations be characterized?
- What representations and algorithms are needed to generate and understand these kinds of descriptions?
- What models and algorithms can be shared with other parts of the system (vision, planning)? How?

1. Discourse Semantics

Our preliminary studies of task related dialogs have revealed some general paradigms for these dialogs. For example, an assembly dialog consists of many subparts of the form

→ GET NEXT PART → PUT IN PLACE → FASTEN.

Each of these subparts may be many sentences long, and each may include subparts. For example FASTEN might include

DETERMINE HOW FASTENED → GET TOOLS [EXTRA PARTS (maybe)] → TIGHTEN.

The most interesting property of this hierarchical structure is that references operate mostly within a subpart. A reference that goes outside the current subpart must be highly specified; e.g., it must indicate

what other subpart the object being referred to was mentioned in. So, for example, inside of GET TOOL there may be such elements as references (e.g., "it," "them") to different kinds of tools and descriptions of differences among tools. However, once the tool has been located, references operate among elements of FASTEN (e.g., what is being fastened). In effect, the discourse pops up a level. Consider the following dialog between the expert and the apprentice:

E: You need an Allen wrench to tighten the screws.

A: What size?

E: 1/16 inch.

A: I can't find one. Where are the wrenches?

E: They're in the top left corner of the toolbox.

A: O.K. Got it. How tight do they have to be?

(The last "they" refers to the screws, not the wrenches.)

To understand this kind of dialog, a structured history of the discourse must be kept. A major thrust of our research will be to characterize this structure and to determine a representation for it and algorithms for using it to aid understanding.

2. Pragmatics and Model Interaction

The semantic component of the speech system will need a detailed model of the devices being worked on in the workstation to understand both what is being said and what the state of the task is. For example, it will need to know what supports (or is supported by) a part, how parts are fastened to each other, and what spatial relations exist between parts. This kind of information is needed to understand statements like, "The washer on top of the spring is worn," or "I removed the belt housing cover."

We will have to define general operators for working on these models. As well as aiding understanding, these operators will form an important communications link between the language and planning components of the system.

"I put the nut on the bolt," and "I put the nut on the table," mean very different things. These different senses of "put" will be expressed by different operators (e.g., "screw-on" and "place-on"). It is these operators that will be communicated to the planning component.

Much of this task-related information will be needed by the vision and planning components of the system. To construct the discourse history, the system will have to know what subsequent states and what problems can follow from a given state. Computing these consequences of a state is really what is meant by planning. Thus, a second major area of research will be determining what task-related information the language part of the system needs and how this will interact with the operators and models in the planning and vision components.

3. Description Generation and Comprehension

Descriptions of parts, tools, and locations are an integral part of expert-apprentice dialogs. In order to communicate effectively with people, the system will have to be able to understand (generate) descriptions that are easily produced (understood) by people. Information from both the preceding dialog and the workstation environments will play an important role.

We first consider the linguistic problems that arise in the construction and comprehension of descriptions. The system will have to be able to use and understand reference and substitution. For example:

E: Find a set of four bolts. Each one has a nut
and three washers attached to it.

A: O.K. I've found four bolts.

E: Put the bolts in and leave them-loose.

A: What should I do with the nut after I put the bolts in place?

Once the system has used "four bolts" it must refer to them as "the (four) bolts" or "them." If it uses "four bolts" (indefinite), the user will be confused; he will probably think the system means another four bolts.

The system will also have to handle ellipses. Consider the following dialog:

E: Have you loosened all the bolts?

A: No, just the front ones.

The last sentence is a fragment. It is equivalent to the statement "I have loosened only the front bolts," although neither "loosening" nor "bolts" appear in it.

Planning problems are also involved in generating good descriptions. Ease of user understanding means the detail of a response will depend on its context (the level of the current task, whether the response is a followup to a previous response or a new topic, the amount of detail the user has required previously, and the kinds of features which distinguish the object being described from others in the current context). The question, "Where on the platform does the pump go?" might be answered, "Near the left side; you can tell its position by where there is no paint." This has obvious advantage over answering, "6.25 inches from the left side and 4.76 inches from the front edge," or worse yet, "At position 6.25, 4.76." Similarly, "Which wrench?" can be answered, "The big one," if there are only two, or "The one with the red handle," if there are many with different colored handles. It is necessary to answer, "The 1/8-inch one," only if there are many and they all look alike.

The system will have to be aware of the user's orientation. Front, back, and side are all relative descriptors. If it is not sure, the system will have to orient the user, e.g., by telling him to face the compressor with the highest part of the compressor away from him. Once it is aware of the user's orientation, the system can give descriptions from the user's point of view.

Ideally, seeing from the user's point of view would be considered in its less literal sense. The system would give descriptions in terms of characteristics the user has already indicated he is aware of. This does not preclude teaching the user new names, tools, or techniques; it only says they will be explained in terms he can understand easily.

It is important to note that people do not usually describe objects in the fewest bits possible. Instead, they seem to intend to enable the listener to locate the object as rapidly as possible. Descriptions involve a narrowing of focus (context). Typically, the initial part of a description includes a coarse physical characterization of the object and an indication of its location. For example, one response to the question, "What's a nut driver?" was, "It looks like a screwdriver and is in the yellow case by the wall." Similarly, "What's the belt housing cover?" was answered, "The vertical piece of metal at the rear of the platform." Locations are also described by narrowing the focus.

E: There is a pair of other nuts.

A: Other nuts?

E: There is a pair of nuts that are sitting on the desk. They are close to a small plate about three inches long.

Functional descriptions also serve to speed up the listener's search. To say that an object is used for some job is to say that it looks like it should perform that job. When there are several objects of the same color and material in approximately the same location, shape

may be the only distinguishing feature. If the object to be found has a complex shape, the simplest and clearest way of describing it may be to tell its function. Consider the following excerpt from an assembly dialog:

E: Now we need to attach the conduit to the motor.
The conduit is the covering around the wires
that you were working with earlier. There is
a small part...um.... Oh brother.

A: Wait a sec.... The conduit is the cover to
the wires?

E: Yes.

A: Oh, I see, there's a part that's supposed to
go over it.

E: Yes.

A: I see. It looks just the right shape too.
Ah-hah!

This kind of description based on narrowing of context and distinguishing features of objects looks very much like what the vision part of the system will be doing except that it is geared toward features that will distinguish an object for a person (instead of for a machine). There is probably a different set of primitives, but the planning algorithms should look a good deal alike. The third major area of research will be determining these primitives and the algorithms needed to decide what the easiest ways of describing an object are. This work will involve the vision and planning components of the system both in sharing algorithms for determining distinguishing features and for determining what the best descriptions are. For example, pointing can often be used, partially or completely, to replace a verbal description. No part of the system can make this decision alone.

D. System Integration

There are major design problems involved with putting together an overall system of this sort. We have already mentioned the desirability of the subsystems sharing model information. In fact, the discussion here has been under separate headings mainly for ease of expression. In reality, the overall system will appear much more homogenized. Natural language and vision tasks will use planning abilities at almost every step, so we can expect a thorough interdependence between parts of the system.

A key feature of the system design will be modularity. The domain of knowledge possessed by the system will consist of far too many items, related in far too many ways, for the system to be constructed in a monolithic fashion. Rather, individual microconsultants will be constructed to deal with separate microworlds (e.g., troubleshooting an electric motor; using a wheel puller). The overall system will assume the responsibility of integrating the efforts of the various microconsultants by dealing with the interactions between them.

Strict adherence to a modular design will also make our system extendible. It is our goal to have the system extendible in the following ways:

- It will be possible to add new factual information about properties of tools, equipment, and such directly to the system without having to alter any of the model maintenance, planning, deduction, or retrieval routines and without having to restructure the existing model information. In particular, it will be possible to teach the system about new equipment (perhaps equipment similar to the compressor).
- It will be possible to add new information about how parts fit together, how to troubleshoot, whatever other knowledge is needed in planning. This new knowledge may consist of more detailed or augmented information about already known equipment or new information about

additional equipment. Often, this kind of knowledge will be in the form of new programs to be added to the complex of old ones, but it will not be necessary to reprogram the systems that use this knowledge.

- There will be some capability for the system to extract new knowledge automatically through dialog with the apprentice or a human teacher and through vision.

We expect additional research problems concerning system integration to surface during the course of the project. For this reason we shall be attempting the construction of a complete system early in the project to provide an opportunity to attack these problems.

IV PROJECT PLAN

A. Phasing

In this section we shall present our plan for achieving the objectives of our project. There is no doubt that the specifications for the demonstration system are challenging. It seems to us that such a system could not be achieved in less than five years, and even then, success will depend on good organization and hard work. Still, we do not think it productive to set our sights lower and plan for a shorter period, and we recommend a five-year project beginning April 1973.

It also seems reasonable to us to divide the project into two consecutive phases. Our reasons for favoring such a division are that

- (1) It gives us a clear mid-project opportunity to review the course we have taken and make major modifications to the plan and our method of approach if warranted.
- (2) It allows us, during the first phase, to concentrate our work on an appropriately scaled-down demonstration to test the feasibility of our approaches.
- (3) It gives a convenient point for ARPA to evaluate our progress.
- (4) It provides a reasonable point at which we can introduce any needed additional computational facilities and reprogramming.

Therefore, we recommend two phases. Phase I, the first three years, will conclude in April 1976 with the demonstration of a scaled-down system. Phase II, the last two years, will be devoted to completing the work needed to demonstrate the final system.

B. The Phase I System

The Phase I demonstration system will be a scaled-down version of the Phase II system. We expect that building it, a much less ambitious system, will teach us which approaches are feasible and which are not.

The task domain for the Phase I system will be the same as that for the final one except that the equipment to be checked out and repaired will consist of a single, simple air compressor, and the system will be designed to work with a single user. The workstation will be largely the same as the Phase II workstation but with fewer tools. Also, we will not be attempting, during Phase I, to achieve real-time operation. That is, in some cases, there may be delays on the order of 2 to 10 minutes between user queries and system responses.

The Phase I system will be demonstrated by April 1976. We will give more details about it as we describe the specific research milestones later in this section.

C. Organization

We can roughly divide the work of the project into four major activities: management, system building, system support, and research support. By management, we mean the obvious activities of monitoring the project plan and continuing its development plus day-to-day running of the project. The system-building activity will involve integrating the results of our research into smoothly running demonstration systems. By system support, we mean the tasks connected with using, maintaining, and improving our systems hardware and software. By research support activity, we mean generating the scientific and technological basis for the demonstration systems.

Our resources will be split roughly as follows:

Management (Leader: Nils Nilsson)	10%
Systems Support (Leader: Daniel Lynch)	25%
Systems Building and Research Support (Leader: Richard Fikes)	65%

Systems building effort will be small at the beginning of each phase and grow larger as the demonstration times approach; the opposite will be true for research support effort.

D. Research Tasks

The research support activity will initially be divided into the following major research tasks:

- Planning and execution monitoring
- Vision research
- Natural language.

1. Planning and Execution Monitoring (Leader: Richard Fikes)

a. Present Activities

Three major present activities are now underway. These are the compressor simulation, the assembly program, and the reconfiguration planner.

--Compressor Simulation (Richard Fikes). A simulation of the actions of the top-level components of a simple air compressor has been written in QLISP.¹³ This simulation makes use of QLISP assert teams (similar to PLANNER antecedent theorems and to QA4 demons) to model the dynamics of a compressor. It is now being extended to more detailed levels, that is, to simulations of the pump and other components. This simulation package will be used mainly in planning for troubleshooting

activities. In developing this simulation, we have created special QLISP constructs for modeling the dynamic states of the compressor (how much air is in the tank, speed of motor, and so forth) and the ways in which these states change.

--Assembly Program (Earl Sacerdoti). Earl Sacerdoti has written a small QLISP program to study the problem of how a computer system should give advice to a human and model his progress in the assembly of components. The program gives hierarchically organized instructions to be executed by the apprentice. If the apprentice informs the system that he cannot carry out an instruction (perhaps he doesn't know how), the system gives more detailed instructions on how to accomplish it. The system updates its model dynamically at the appropriate level of detail. This system is the basis for a program called the Procedural Net System now being developed.

--Reconfiguration Planner (Richard Fikes). A planning system is being designed that will be able to give an apprentice instructions on how to reconfigure (e.g., assemble, disassemble, retrofit, modify) an air compressor from any arbitrary state to any other. This program will be basic to any task requiring physical operations with the compressor. This planning system will probably merge with Sacerdoti's Procedural Net System.

To date, we have worked out some preliminary ideas for modeling information about how compressor parts are to be connected and disconnected and for representing an arbitrary configuration of the compressor.

b. Plans for the Future

--1974. During the coming year, we expect to have a system running that can plan instructions for reconfiguration of the top-level compressor components (e.g., pump, motor, fan belt, tank, pressure switch). These instructions will be given at any of several levels of detail depending on responses from the user. We also plan to investigate how information about failures can be used to guide planning.

As a part of this system we will develop various QLISP associative deductive retrieval routines for querying the data base. These will be based on the QA4 routines already written.⁸ They will be used by the planning system itself, but they will also be available for question answering. Thus, by the end of 1974, we will be able to demonstrate a system having the ability to answer questions about compressor parts, connectivities, properties of parts, and so forth.

During 1974, we will be developing a store of information about tools and their uses largely in the form of routines that are used to plan actions using tools.

We will continue work on dynamic model updating by making use of environmental information. During 1974, this information will be supplied primarily through symbolic (rather than visual) interaction with the apprentice. This updating ability will be used during late 1974 and later by plan monitoring routines.

Finally, we will develop a representation for plans that will also allow us to model the progress and history of the task at hand. This will be important for monitoring and replanning activities.

--1975. By the end of 1975, we will have a multilevel reconfiguration planner running (i.e., one that can reconfigure the parts of the pump and other components as well as those at the top level). Our

model information on which this planner is based will now contain some information about the user, for example, information about his level of expertise. We will also begin to put in and use some specialized information about troubleshooting. A beginning planning system for troubleshooting will be in operation by the end of 1975. This will be improved during 1976 and integrated with the reconfiguration planner for the Phase I demonstration.

Earl Sacerdoti is going to be working on a Ph.D. dissertation in this general research area (through the Computer Science Department at Stanford University), and his dissertation ought to be finished during 1975. In addition we will continue to encourage the participation of other Stanford students.

The planning capabilities that have been achieved by these research tasks by the end of 1975 define the planning capabilities of the Phase I (April 1976) demonstration system. We think that it is not productive to generate detailed plans now for the course of problem-solving research during Phase II. These will be developed as the work in Phase I proceeds.

2. Vision Research (Leader: J. M. Tenenbaum)

a. Present Activities

We are now working in three mutually complementary vision research areas that will begin to merge in 1974. First, we have constructed an interactive scene analysis system that is used for synthesizing a knowledge base for analyzing scenes.¹⁴ Given a particular domain, the researcher can use the system to experiment with the utility of various perceptual operators (such as color, texture, brightness, height, orientation) for recognizing objects in the scene. Through this interactive process, specialized scene analysis programs can be quickly constructed.

(The interactive scene analysis system has resulted from joint sponsorship by ARPA, ONR, and NASA.)

A second activity concerns developing a highly goal-directed system for answering questions about a scene. Given a question, this system will automatically organize the appropriate perceptual operators (based on the semantics of the domain) to locate those objects in the scene that are relevant to the answer. For example, in response to a question such as "Is the compressor plugged in?", the system would first use the distinguishing features of a wall power socket (perhaps the color of the socket plate) and of the compressor power plug (dull black) to locate these objects. If the power socket area does not contain the plug, the system might answer no; or it might choose to locate the plug itself, knowing the distinguishing features of the power cord, to provide confirmation. At the present time, we are completing the implementation of this system. It is based largely on the doctoral research of Thomas Garvey, who expects to obtain a Ph.D. degree from Stanford University during 1974.

Garvey's system will typically be adequate for answering many questions about scenes; however, it has obvious limitations as a general perceptual strategy. In complex environments, it may be impossible to assign unique interpretations to parts of a scene based only on local perceptual attributes taken out of context; different objects can have similar appearance, while objects belonging to the same functional class can have strikingly different appearances (e.g., wrenches). Occlusion and image imperfections contribute to variabilities in appearance. In these cases, nothing will suffice short of a complete analysis of the scene.

As a third activity therefore, we are currently developing a semantics-based scene-understanding paradigm, in which interpretations

that are ambiguous (based on local perceptual attributes of regions) are ruled out using contextual constraints in order to arrive at a consistent and meaningful interpretation of the whole scene. This system is intended to assimilate a wide variety of independent facts about a pictorial domain and to produce an interpretation whose quality should improve with the amount of available knowledge.

In addition to these activities, we are continuing the development of a time-of-flight laser range finder¹⁵ that will provide depth data for picture fields on the order of 100-by-100 points.

b. Plans for the Future

Our future work in vision has been arranged into a set of tasks spanning the next five years. These tasks are largely based on our present work and a carefully thought-out vision research program^{16, 17} that now focuses on the computer-based consultant project.

--1974. Effort during 1974 will consist of

- Demonstration of a working laser and associated software to acquire, display, transform, and calibrate 3-dimensional image data.
- Identification of parts on an assembled compressor in known position and orientation designated by the system's laser pointer, a hand held LED, or a display cursor. (Identification will be accomplished without explicit use of vision by transforming the part designation to a known ray in space and computing the intersection of that ray with a projected geometric model of the compressor.)

--1975. Effort during 1975 will include the following

tasks:

- Determine the orientation of isolated compressor parts whose identity is known.
- Generate automatically "distinguishing feature" strategies for finding objects in room scenes from declarative descriptions.
- Analyze workstation scenes using global constraint satisfaction methods.
- Accomplish visual recognition of isolated parts and tools based on geometric models and 3-dimensional image data.
- Accomplish identification of parts on an assembled compressor designated by a man pointing with his finger.

--1976. Effort during 1976 will include the following

tasks:

- Generate distinguishing feature strategies for locating a specified part or tool anywhere in the work station.
- Identify all parts in a subassembly and be able to determine interconnections. Check part alignment and so forth. (possibly utilizing constraint satisfaction methods).

--1977. Effort during 1977 will include the development of man-machine pictorial communication. The system will be able to recognize parts from human verbal descriptions and, conversely, to describe parts to an apprentice making effective use of the user's contextual referents. Descriptions will be in list format at first, but, by the end of the year, they will be in natural language at the competence level of available parsers.

--1978. Effort during 1978 will include development of an integrated (man-machine) information gathering capability with sufficient problem-solving expertise to provide the range of services needed by the Phase II (1978) demonstration system.

The visual abilities of the Phase I (April 1976) demonstration system are defined by the vision tasks through 1975. In summary, the Phase I system will be able to recognize isolated parts and tools, and it will be able to identify the parts on an assembled compressor as they are designated by a man pointing with his finger.

3. Natural-Language Investigation (Leader: Barbara Deutsch)

a. Present Activities

We are beginning to investigate the syntax, semantics, and vocabulary of the workstation domain. Various protocols have been tape recorded of dialogs between a human expert and human apprentice while the apprentice was assembling and disassembling the compressor. From these, we will estimate the requirements for vocabulary, syntax, and semantics for a continuous speech-understanding system for this domain.

b. Plans for the Future

--1974. So that we can have natural-language input as soon as possible, we will adapt the SRI Speech-Understanding System Parser¹⁸ to accept text input from the workstation domain. For the five-year system, both the syntax and semantics of the system will have to be extended. We expect to take advantage of developments in the speech-understanding project. A major goal in the first year of the project will be to determine the syntactic constructions used in goal-directed person-machine communication. We will begin to extend the grammar in this direction.

Our focus in 1974 will be on the interface between natural language input/output and the problem-solving and planning models. We will investigate the kinds of information these processes need to share and how they will interact. For example, all of these processes will need "models" of the tools and parts in the workstation domain. Natural language processes need these models in order to interpret the input in terms of basic operations in the domain. Planning processes need them to decide which tools and parts to use at which times.

We will also be investigating how dialogue and task-related information can be used to aid understanding in planning, problem solving, and natural language input/output. We will begin modeling discourse structure and testing representations and algorithms for incorporating it in the natural-language understanding system.

Also during this year, we expect to use a commercially available, isolated-word recognizer as a simple input device. Such a device can be used for inputting such phrases as "What?" "I don't understand," "Okay," and "More detail please."

As regards natural-language output from the system to the apprentice, we shall postpone developing any sophisticated system until 1975. During 1974, the system's responses will be selected from canned sentences that will be typed out and, toward the end of the year, spoken on a VOTRAX speech synthesizer.

--1975. During 1975, our text input system should be working and understanding simple sentences from the apprentice. Modifications and extensions to the grammar will continue. We will be testing and modifying our semantic models for the tools and parts in the workstation domain. We will continue the work on modeling discourse structure. Emphasis will be on how natural language can share representations and algorithms with other parts of the system.

Work will also continue with the isolated-word recognizer, and we shall begin work on automatic synthesis of sentences. Our plans for beyond 1975 are not worked out in detail, since much of what we will do depends on the direction of current speech-understanding research. We do propose to postpone any acoustic level work until 1976. However, we will be able to take advantage of the acoustic work done in the speech-understanding project.

Barbara Deutsch expects to be doing Ph.D. research through the University of California at Berkeley in this general area.

The Phase I system will have only a limited natural-language capability, and this will probably be mainly through the isolated-word recognizer and a text input system.

Appendix A

DESCRIPTION OF A SIMPLE AIR COMPRESSOR

Appendix A

DESCRIPTION OF A SIMPLE AIR COMPRESSOR

The air compressor we have been working with is Sears Model 17209-1/2 HP Compressor. (It was illustrated in Figure 2.)

The essential elements of this unit are

- An electric motor (drives the air pump).
- A pump (takes air from atmosphere and compresses it into tank).
- A tank (stores compressed air).
- A pressure switch (starts and stops the motor at pre-determined pressures).
- A relief valve (releases "back" pressure on pump when the motor stops).
- A check valve (holds air in the tank when compressor stops and allows the relief valve to function).

We quote from the Sears operating instructions to describe the major features:

PUMPING. In action, when the piston of the pump goes down, air rushes in through the inlet valve. As the piston starts up the inlet valve closes and the air in the cylinder is compressed until the pressure in the cylinder slightly exceeds the pressure in the piping (aftercooler) connected to the check valve. This excess pressure causes the outlet valve to open, allowing the air in the cylinder to be forced out through the outlet valve into the piping. As the piston starts down, the outlet valve closes, preventing air in the aftercooler from rushing back into the cylinder and the cycle starts over again. When the pressure in the aftercooler exceeds that in the tank, the check valve opens and allows air to pass into the tank.

STOPPING. The compressor continues to pump until the pressure in the tank reaches the point at which the automatic switch is set to cut off. The pressure in the tank acts on the diaphragm of the automatic switch (through piping) and when the cut-out point is reached, this diaphragm distends and, through levers, opens switch contacts stopping the motor, and simultaneously opening the relief valve. The air trapped between the pump and the check valve is thus allowed to escape through the relief valve. The air in the tank cannot, of course, escape because the check valve does not allow the air in the tank to flow back.

STARTING. When the pressure in the tank is decreased by using the air, and finally reaches the "cut-in point" of the switch, the reduced pressure on the diaphragm allows the pressure spring to cause the lever to close the switch contacts and relief valve. This starts the motor and the compressor again begins to pump. There being no pressure in the aftercooler (it having escaped when the relief valve was opened) the motor can start under no pumping load and get up to speed during the time required to raise pressure in the aftercooler.

From the foregoing and from an examination of the diagram, it will be seen that not only must the electric motor and air compressor itself be kept in proper working condition, but that the check valve and relief valve must work properly. If the check valve leaks, then the air stored in the tank will escape through the relief valve when the compressor is not running because the relief valve is open all the time the switch contacts are open.


Also, if the relief valve doesn't close properly when the compressor is running, the air (or most of it) pumped by the compressor will not go into the tank at all, but will pass out through the leaky valve. Or, if the relief valve does not open properly then the back pressure can not be relieved and the motor must start the compressor against pressure, which consumes extra power and may damage the motor.

The major parts of the compressor are listed in Figure A-1 (taken from the Sears manual).

The parts list for the pump unit is given in Figure A-2.

Model Numbers

(17199) 102.17200 102.17220 (17218) 102.17021
(17209) 102.17210 102.17001 102.17046
102.17011 102.17041




AUTOMATIC SWITCH CONTROL - STORAGE TANK TYPE

PARTS LIST

Reference Number	Part Number	NAME OF PART	Reference Number	Part Number	NAME OF PART
1	690	Aftercooler Elbow 1/8" Flare Tube x 1/4" Pipe	13	12191	Pressure Switch for (17209) 102.17210, 102.17220 (17218)
2	-	Pump - See Figure 9	13	13858	102.17011, 102.17021 and
3	37055	Aftercooler for 102.17001, 102.17011, 102.17021, 102.17046 and 102.17041	13	32473	102.17041
3	37057	Aftercooler for (17199) 102.17200, (17209) 102.17210, and 102.17220 (17218)	14	-	Pressure Switch for 102.17046
4	98556	Elbow Connector-1/4" Tube x 1/8" Pipe Thread	15	18591	Pressure Switch for (17199) 102.17200
5	15681	Check Valve-See Figure 4 for Parts	16	18550	Pressure Switch Relief Valve. When ordering
6	12776	Nameplate	16	18544	Relief Valve give Pressure Switch Nameplate
7a	9609	Oil Drain Nipple-1/8" x 5" Long	17	-	Data. (Ref. No. 13)
7b	14292	Oil Drain Cap-1/8"	18	45053	Check Valve Relief Tube (also available
8	43779	Tank and Platform Assembly for (17199) 102.17200, (17209) 102.17210, 102.17220 (17218)	18	45058	locally 1/4" aluminum or copper tube 31 1/2 long.
8	43780	Tank and Platform Assembly to Massachusetts Specifications for (17199) 102.17200, (17209) 102.17210, and 102.17220 (17218)	18	45055	Safety Valve for 102.17001, 102.17011,
8	43781	Tank and Platform Assembly 102.17001	18	45059	102.17021, 102.17046 and 102.17041
8	43782	Tank and Platform Assembly for 102.17021, 102.17046 and 102.17041	18	30545	Massachusetts Safety Valve for (17199) 102.17200,
8	43783	Tank and Platform Assembly to Massachusetts Specifications for 102.17021, 102.17046 and 102.17041	19	17469	(17209) 102.17210, 102.17220 (17218), 102.17001,
9	16614	Manifold for mounting Pressure Switch, etc. for 102.17021, 102.17046 and 102.17041. Use for all Mass. Specifications	19	17470	102.17011, 102.17021, 102.17046 and 102.17041
9	43724	Manifold for mounting Pressure Switch and Gauge, (Safety Valve and Discharge Valve included) for (17199) 102.17200, (17209) 102.17210, 102.17220 (17218) and 102.17001. Not for Massachusetts Specifications.	19	18356	Motor or Gas Engine-See Motor Nameplate and
10	12571	Pressure Gauge for 102.17001, 102.17011, 102.17021, 102.17046 and 102.17041	19	18481	Gas Engine Part Booklet for characteristics;
10	15793	Pressure Gauge for (17199) 102.17200, (17209) 102.17210, and 102.17220 (17218)	19	32622	when writing regarding motor or engine service
10	17088	Pressure Gauge for State of Massachusetts	19	96570	include complete nameplate data.
11	95939	Discharge Valve for 102.17001, 102.17011, 102.17021, 102.17046 and 102.17041. Also used on Mass. Specification models (17199) 102.17200, (17209) 102.17210 and 102.17220 (17218)	Not Shown	96800	Motor Pulley for (17199) 102.17200
12		BX conduit-Available locally (See Installation Instructions for wire size)	Not Shown	13231	Motor Pulley for (17209) 102.17210 and 102.17021
			Not Shown	43429	Motor Pulley for 102.17001 and 102.17011
			Not Shown	43430	Motor Pulley for 102.17220 (17218) and 102.17041
					Motor Pulley for 102.17046
					Belt for 102.17001 (1/2" wide
					x 44" long)
					Belt for 102.17220 (17218), 102.17021 and 102.17041
					(1/2" wide x 47" long)
					Belt for (17209) 102.17210 (1/2" wide x 46" long)
					Belt for 102.17011 and 102.17046 (1/2" wide
					x 45" long)
					Belt for (17199) 102.17200 - 1/2" wide x 43" long
					Aftercooler Vent Coco-Gas Engine Only
					Tank Drain Valve
					Tank Drain Valve Only for State of Mass.
					Cord and Plug Assembly for 102.17301,
					102.17400 and 102.17410
					Adapter Ground Plug for 102.17304
					102.17400 and 102.17410

SA-1530-56

FIGURE A-1 PARTS LIST FOR AIR COMPRESSOR



Reference Number	Part Number	NAME OF PART	Quan. Used	Reference Number	Part Number	NAME OF PART	Quan. Used
1	96472	Head Bolt	2	30	21614	Breather Valve Spacer	1
2	45007	Cylinder Head Cover	1	31	38941	Breather Valve Plate	1
3	45020	Filter Screen	1	32	38944	Breather Valve	1
4	45021	Filter	1	33	45003	End Cover Assembly-Includes Ref. No. 27	1
5	45024	Outlet Valve Stop	1	34**	45025	Pulley for 101 TV	1
6	45014	Copper Washer	2	35	98627	Pulley Set Screw- $\frac{7}{16}$ -18 x $\frac{1}{2}$ " Socket Head	2
7	45011	Outlet Valve Spring	1	36	18060	End Cover Bolt- $\frac{1}{4}$ -20 x $\frac{1}{2}$ "	4
8	45013	Valve	2	37	9522	Lockwasher- $\frac{1}{4}$ "	2
9	45016	Outlet Valve Seat	1	38	11109	Hex Head Bolts- $\frac{1}{4}$ -20 x 1"	2
10	45019	Outlet Valve Retainer	1	39	9522	Lockwashers- $\frac{1}{4}$ "	6
11	45017	Inlet Valve Seat	1	40	980	Hex Head Bolts- $\frac{1}{4}$ -20 x $\frac{1}{4}$ " long	6
12	45012	Inlet Valve Spring	1	41	45006	Base Plate	1
13	45018	Inlet Valve Retainer	1	42	45009	Base Plate Gasket	1
14	18061	Head Bolt	2	43	9349	Pipe Plug $\frac{1}{4}$ "-Furnished with 102,17500 only	1
15	45005	Cylinder Head	1	44	9348	Pipe Plug $\frac{1}{4}$ "	1
16	45008	Cylinder Head Gasket	1	45	9609	Pipe Nipple- $\frac{1}{8}$ " x 5"-Furnished with compressor only	1
17	31901	Piston Rings	2	46	14292	Pipe Cap- $\frac{1}{4}$ "-Furnished with compressor only	1
18	17083	Oil Ring	1	47	45027	Crankcase Plug	1
19	38932	Piston	1	Not Shown	43797	Parts List	.
20	12992	Wrist Pin	1	**	45050	Pulley for 100 TV	1
21	7890	Wrist Pin Retainer Washer	2	*	45051	Crankshaft for 100 TV	1
22	45001	Connecting Rod Assembly	1	Not Shown	45060	Head Plate Spacer	2
23	45004	Crankcase Assembly-Includes Ref. No. 27	1	Not Shown	45244	Set-Gaskets	1
24	45023	Thrust Washer	2	Not Shown	45237	Set-Rings	1
25*	45002	Crankshaft for 101 TV	1	Not Shown	45245	Set-Valve & Spring	1
26	9397	Woodruff Key No. 9	1				
27	45026	Bronze Bushing	1				
28	45010	End Cover Gasket	1				
29	38951	Breather Valve Screw-8-32 x $\frac{3}{16}$ "	1				

SA-1530-57

FIGURE A-2 PARTS LIST FOR THE PUMP UNIT—MODEL NUMBER 102:17500 (17501)

Appendix B

ANALYSIS OF AIR COMPRESSOR CHECKOUT TASK

Appendix B

ANALYSIS OF AIR COMPRESSOR CHECKOUT TASK

At the top level, the system will provide guidance by asking the apprentice to perform each of the steps described below.

1. Step 1

Step 1 in the checkout procedure is to visually inspect the entire compressor, making sure that none of the components are missing or have obvious external damage (e.g., frayed wire, broken belt, split tubing). This might be a reasonable automatic task for the visual perception part of the system, perhaps with help from the apprentice to point sensors (e.g., camera) at various components. To the extent that the apprentice would be doing the actual inspection, he would have to be provided with displays or other descriptive information about the location and appearance of each component to be inspected. When a component is found to be missing (e.g., the belt, the motor), then the system must be able to tell the apprentice how to install a replacement component. When a component is found to be damaged, then the system must be able to describe to the person a step-by-step procedure for either repairing the component or for removing it and replacing it. These assembly, disassembly, and repair procedures should be able to work at varying levels of detail depending on the person's skills. They also should include information as to which tools to use and how to use them.

2. Step 2

Step 2 in the checkout procedure is to make sure the compressor is in a level position with all tank feet firmly supported.

3. Step 3

Step 3 is to make sure all bolts and connections are tight (e.g., those for motor and pump mounts, tubing and gauge connections). This step should probably be presented as a checklist of bolts and connections to test. If the user requires it, the presentation will have to include information as to the location of each bolt or connection. The system must provide some information as to what is meant by "tight" in each case (e.g., finger tight, mild-pressure-with-a-wrench tight).

4. Step 4

Step 4 is to check alignment of the motor and pump pulleys. This involves removing the protective covering and using Allen wrenches to loosen and tighten the pulleys. The instructions for this step should indicate to the apprentice how to test for alignment of the pulleys (e.g., sight along the belt, check if belt is binding).

5. Step 5

Step 5 is to attempt to turn the pump pulley by hand. Turning the pump pulley should cause the belt to move and the pump axle, motor pulley, and motor axle to turn. The pulleys may need tightening; the belt tension may have to be adjusted (by loosening the motor mount bolts and moving the motor); and either the pump or the motor may be "frozen" so that they will not turn. In the case of a frozen motor or pump, the system could suggest replacement of the motor or pump, or it could move the apprentice into a repair mode. Note that the system should coordinate

Steps 4 and 5 so that, if the protective covering was removed in Step 4, then Step 5 should be done before it is replaced, since replacement of the motor or pump would require removal of the covering. Also, the system should realize that, if in Step 4 the pulleys were realigned, then they have been tightened and Step 5 need not concern itself with checking for loose pulleys.

6. Step 6

Step 6 is to examine felt pad in the intake filter. If dirty, replace the felt. This step involves locating the intake filter for the apprentice, teaching him how to remove the felt cover, and indicating in some way how dirty is "dirty". (Maybe the apprentice removes the felt pad and the system decides via the TV camera).

7. Step 7

Step 7 is to check oil level and dirtiness in pump. This step may involve detailed instructions on how and where to check the oil, how dirty, again, must the oil be to need changing, how to change the oil, what kind of oil to use, and so forth.

8. Step 8

Step 8 is to note tank pressure-gauge reading, then open tank valve to allow all pressure to escape from tank, and close tank valve. This step prepares the compressor for running (in the next step) and provides initial checks on the tank valve and the tank pressure gauge. If the tank valve will not open, then the system should instruct the apprentice on how to remove and repair or replace the tank valve. The tank pressure gauge may be assumed to be faulty and should be replaced or repaired if there was obviously pressure in the tank initially and

the gauge read zero or if the gauge does not read zero after the pressure has been released from the tank.

9. Step 9

Step 9 is to plug in power cord. Motor and pump should begin operating, and the tank pressure should begin to rise. The apprentice should be told to check for unusual occurrences such as sparks flying, motor smoking, obvious air leakages, and so forth, and to unplug cord when any are observed. The system should have troubleshooting procedures for some reasonable collection of such symptoms. Given that no such problems exist, the system should indicate the rate at which tank pressure should be rising, e.g., tank pressure at about 25 psi after one minute, 45 psi after two minutes. The system should indicate the expected behavior--namely, motor pump should shut off when tank reaches 125 psi (in about 7.5 minutes); a few seconds after pump shutoff, a brief hiss (about 1 second) will be heard from the check valve as the pressure in the aftercooler is released; tank pressure will slowly drop about five psi during the minute after pump shutoff because of cooling of the air in the tank; tank pressure should then remain at 120 psi. The apprentice should be told not to let motor and pump run longer than, say, eight minutes under any circumstances. The system should be able to troubleshoot failure of the motor and pump to shut off, considering such possibilities as faulty pressure switch, belt slippage, worn rings in pump, air leakage, worn brushes in motor. If the motor and pump shut off at some tank pressure reading other than 125 psi, the system should consider asking the apprentice to check for a faulty gauge and then be prepared to adjust the pressure switch so that shut-off will occur when desired. If the tank pressure does not remain steady at 120 psi, then

the system should direct the apprentice to look for leaks at likely spots (e.g., tank valve, check valve, after-cooler connections) and to correct them when found.

10. Step 10

Step 10 is to open the drain valve for a few seconds to clear out any moisture and close the drain valve. The system should indicate such things as location of drain valve, what tool to use if the valve cannot be opened by hand.

11. Step 11

Step 11 is to open the tank valve and let air escape from the tank until the pump and motor restart and close the tank valve. The motor and pump should restart when the tank pressure reaches 95 psi and should remain on until the tank pressure builds up to 125 psi (about 2.5 minutes). At that time, as in Step 9 above, the motor and pump should stop, pressure should be released from the aftercooler, and the tank pressure should hold steady at 120 psi. If the motor and pump do not start at 95 psi and stop at 125 psi, then the system should direct troubleshooting, repair, and/or pressure switch adjustments to correct the situation as in Step 9 above.

Successful completion of these steps would allow the apprentice to indicate that the compressor is checked out and ready for use.

REFERENCES

1. G. Goldberg, M. C. Pease and A. Waksman, "A Program for a Management Support System," Technical Report 6, Contract N00014-71-C-0210, Stanford Research Institute, Menlo Park, California (November 1973).
2. A. Newell et al., Speech Understanding Systems: Final Report of a Study Group (published for Artificial Intelligence by North-Holland/American Elsevier, Amsterdam/New York, 1973).
3. R. E. Fikes and N. J. Nilsson, "STRIPS: A New Approach to the Application of Theorem Proving to Problem Solving," Artificial Intelligence, Vol. 2, No. 3/4, pp. 189-208 (Winter 1971).
4. R. E. Fikes, P. E. Hart, and N. J. Nilsson, "Learning and Executing Generalized Robot Plans," Artificial Intelligence, Vol. 3, No. 4, pp. 251-288 (Winter 1972).
5. E. Sacerdoti, "Planning in a Hierarchy of Abstraction Spaces," Advance Papers of the Third International Joint Conference on Artificial Intelligence, Stanford University, Stanford, California, August 1973.
6. J. F. Rulifson, "QA4: A Procedural Calculus for Intuitive Reasoning," Ph.D. thesis, Computer Science Department, Stanford University, Stanford, California (November 1972).
7. J. A. Derksen, J. F. Rulifson, and R. J. Waldinger, "The QA4 Language Applied to Robot Planning," AFIPS Conference Proceedings, Fall Joint Computer Conference, Vol. 41, Part II, pp. 1181-1192 (AFIPS Press, Montvale, New Jersey, 1972).
8. B. Raphael, R. Fikes and R. Waldinger, "Research in Advanced Formal Theorem-Proving Techniques," Final Report, SRI Project 8721, NASA Contract NASW-2086, Stanford Research Institute, Menlo Park, California (August 1973).
9. T. A. Winograd, "Procedures as a Representation for Data in a Computer Program for Understanding Natural Language," Ph.D. thesis, Department of Mathematics, Massachusetts Institute of Technology, Cambridge, Massachusetts (1971).

10. G. J. Sussman, "A Computational Model of Skill Acquisition," Technical Report TR-297, MIT Artificial Intelligence Laboratory, Cambridge, Massachusetts (August 1973).
11. R. E. Fikes, P. E. Hart, and N. J. Nilsson, "New Directions in Robot Problem Solving," in Machine Intelligence 7, B. Meltzer and D. Michie, Editors (Edinburgh University Press, Scotland, November 1972).
12. N. J. Nilsson, "A Hierarchical Robot Planning and Execution System," Technical Note 76, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (April 1973).
13. R. Reboh and E. Sacerdoti, "A Preliminary QLISP Manual," Technical Note 81, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (August 1973).
14. J. M. Tenenbaum et al., "A System for Interactive Scene Analysis," Annual Technical Report, SRI Project 1187, ONR Contract N00014-71-C-0294, Stanford Research Institute, Menlo Park, California (in preparation).
15. P. E. Hart et al., "Artificial Intelligence--Research and Applications," Annual Technical Report, SRI Project 1530, ARPA Order No. 1943, Contract DAHCO4-72-C-0008, Stanford Research Institute, Menlo Park, California (December 1972).
16. J. M. Tenenbaum, "On Locating Objects by their Distinguishing Features in Multisensory Images," Technical Note 84, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (September 1973).
17. B. Raphael et al., "Research and Application--Artificial Intelligence," Final Report, SRI Project 8973, NASA Contract NASW-2164, Stanford Research Institute, Menlo Park, California (December 1971).
18. W. Paxton and A. Robinson, "A Parser for a Speech Understanding System," Advance Papers of Third International Conference on Artificial Intelligence, pp. 216-222, Stanford University, Stanford, California, August 1973.