AN APPROACH TO ACQUIRING AND APPLYING KNOWLEDGE

Norman Haas and Gary G. Hendrix
SRI International
333 Ravenswood Avenue
Menlo Park, California 94025

ABSTRACT

The problem addressed in this paper is how to enable a computer system to acquire facts about new domains from tutors who are experts in their respective fields, but who have little or no training in computer science. The information to be acquired is that needed to support question-answering activities. The basic acquisition approach is "learning by being told." We have been especially interested in exploring the notion of simultaneously learning not only new concepts, but also the linguistic constructions used to express those concepts. As a research vehicle we have developed a system that is preprogrammed with deductive algorithms and a fixed set of syntactic/semantic rules covering a small subset of English. It has been endowed with sufficient seed concepts and seed vocabulary to support effective tutorial interaction. Furthermore, the system is capable of learning new concepts and vocabulary, and can apply its acquired knowledge in a prescribed range of problem-solving situations.

I INTRODUCTION

Virtually any nontrivial artificial intelligence (AI) system requires a large body of machine-usable knowledge about its domain of application.* Construction of a knowledge base is currently a tedious and time-consuming operation that must be performed by people familiar with knowledge representation techniques. The problem addressed in this paper is how to enable computer systems to acquire sets of facts about totally new domains from tutors who are experts in their own fields, but have little or no training in computer science. In an attempt to find a practical solution to this problem, we have developed a pilot system for knowledge acquisition, which, along with several related research issues, is discussed below.

The kinds of information we are most interested in acquiring are those needed to support what have been called "question-answering" or "fact-retrieval" systems. In particular, our interest is in collecting and organizing relatively large aggregations of individual facts about new domains, rather than in acquiring rules for judgmental reasoning. This is in contrast to previous work on such systems as those of Davis [1] and Dietterich and Michalski [2], that treat knowledge not so much as a collection of facts, but as a set of instructions for controlling the behavior of an engine.

The type of acquisition process we are exploring is "learning by being told," in contrast to the idea of "learning by example." It is this latter concept which has formed the basis of research by other investigators in this area, such as Winston [11] and Mitchell [8].

Our interest in knowledge acquisition is motivated by the desire to create computer-based systems that can aid their users in managing information. The core idea is that of a system that can talk to a user about his problems and subsequently apply other types of software to meet his needs. Such software would include data base management systems, report generators, planners, simulators, statistical packages, and the like. Interactive dialogs in natural language appear the most convenient means for obtaining most of the application-specific knowledge needed by such intelligent systems.

II KNOWLEDGE ACQUISITION THROUGH ENGLISH DIALOGS

Systems that acquire knowledge about new domains through natural-language dialogs must have two kinds of special capabilities. First, they must be capable of simultaneously learning both new concepts and the linguistic constructions used to express those concepts. (This need for simultaneous acquisition of concepts and language reflects the integral connection between language and reasoning.) Second, such systems must support interactive, mixed-initiative dialogs. Because a tutor may provide new knowledge in an incremental

* This research was supported by the Defense Advance Research Projects Agency under contract N00039-79-C-0118 with the Naval Electronic Systems Command. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency of the United States Government.

and incomplete manner, the system must keep track of what it has already been told so that it can deduce the existence of missing information and explicitly ask the tutor to supply it.

We are exploring the feasibility of such ideas by developing a series of Knowledge-Learning and -Using Systems (KLAUS). A KLAUS is an interactive computer system that possesses a basic knowledge of the English language, is capable of learning the concepts and vocabulary of new subject domains, and has sufficient expertise to apply its acquired knowledge effectively in problem-solving situations.

III RESEARCH ISSUES FOR KNOWLEDGE ACQUISITION

To create systems capable of acquiring knowledge through tutorial dialogs in English, several fundamental research problems must be resolved:

A powerful natural-language processing capability is required. Although much progress has been made in recent years, previous work has assumed a complete knowledge base. Knowledge-acquisition dialogs require several adaptations and extensions.

Seed concepts and seed vocabulary must be identified for inclusion in a core system. It is not at all obvious which words and concepts will be most useful in helping tutors describe the concepts of new domains.

A structure for lexical entries must be specified so that the system can acquire new lexical information. Because such information provides a key link between surface linguistic form and underlying meaning, structural specification is a very challenging task for certain categories of words, particularly verbs.

The linguistic constructions that people use in introducing new concepts must be identified and analyzed so they can be interpreted correctly by the natural-language processing system. Such constructions range from simple syntactic patterns to complex uses of analogy.

A flexible scheme of knowledge representation is necessary. The representation must have general expressive power, since it may be applied to many different domains and must support the addition of new information. It should include inherent features that can aid in organizing knowledge and in supporting the incremental acquisition of knowledge.

An efficient problem-solving capability is needed to answer questions and to draw inferences for integrating newly acquired information. This capability must be based on general principles, because no application-specific problem-solving procedures will be included in the system. (How to acquire such procedures is a separate and interesting research question.)

A methodology is needed for integrating new concepts into the system's knowledge base. Because tutors will often provide only partial descriptions of new concepts, methods have to be devised for ascertaining what additional facts must be sought from the tutor to insure proper linkage between the new concepts and those previously acquired.

A set of readily understandable questions is needed for eliciting information from tutors. The length and number of questions should be minimized to impose as small a burden on tutors as possible.

These problems must be dealt with in an integrated manner, balancing the requirements of one facet of the system against those of others. Our initial attempts to cope with this complex of issues are presented below.

IV THE NANOKLAUS SYSTEM

A. Overview

We have recently developed and tested a pilot KLAUS, called NANOKLAUS. The principal components of NANOKLAUS are a natural-language processing module based on LIFER [4], a formal deduction module that operates on a data base of well-formed formulas (wffs) in a many-sorted first-order logic, and a number of support procedures that aid in acquiring knowledge about new subject domains and maintaining the data base. First-order logic was chosen as the basis for our knowledge representation scheme because of its generality and because of the computational soundness and power of problem-solving systems that use it.

B. Seed Concepts and Vocabulary

NANOKLAUS comes preprogrammed with a fixed set of syntactic and semantic rules covering a small subset of English. It also comes with seed concepts and a seed vocabulary, which are to be extended during tuition. For example, the system knows about such things as PHYSICAL OBJECTS, PERSONS, MEASURES, and the like. It has a knowledge of the basic function words of English, as well as of such words as "unit," "kind," and "plural" that are used frequently in stating definitions of new words and concepts. These seed concepts allow the untutored NANOKLAUS to engage in conversations like the following:*

73> <u>A length is a measure</u>
You're saying that anything that is a
LENGTH is also a MEASURE.

----
* The tutor's inputs are underlined in these examples. All other text is produced by NANOKLAUS.

OK, now I have assimilated LENGTH.
74> A foot is a unit of length
FOOTS are the first units of LENGTH you've defined. Now I have assimilated FOOT.
75> The plural of foot is feet
FEET - got it. Thanks.
76> A meter is a unit of length
How is it related to FOOT? >> There are 3.3 feet in a meter
Now I have assimilated METER.
77> A physical object has a length
So PHYSICAL OBJECTS have LENGTHS.

Exactly what to use as seed concepts is problematical for a system that must bootstrap its entry into a new domain. Most of the seed concepts in NANOKLAUS are classes of THINGs and RELATIONs. They have been included in the system either simply to avoid forcing tutors to relate everything to the most general concept (THING), or because they have a special status in English. For example, because pronominal usage depends on gender, the class MALE-BEING is defined and associated with the pronoun "he." One might consider defining as seed concepts a set of primitives, to which all other concepts must be reduced, but such a reductionist approach is probably unworkable [3].

NANOKLAUS uses five principles of knowledge organization to integrate new knowledge: (1) there are things; (2) there are subclasses of things (i.e., things can be subclassified); (3) there are relations among things; (4) there are subclasses of relations; (5) some of the relations are functions. The concepts of uniqueness and equality also play important roles. NANOKLAUS is not programmed to hold explicit conversations about these concepts, but rather to use them in its internal operations.

C. The Natural-Language Component

The natural-language component of NANOKLAUS uses a pragmatic grammar in the style of LADDER [6]. Although most of the linguistic processing performed by the system follows fairly standard practice, the pragmatic grammar is distinguished by its explicit identification of a number of syntactic structures used principally to define new concepts. As an oversimplified example, NANOKLAUS might be thought of as looking for the syntactic pattern

<S> => <A> <NEW-WORD> <BE> <A> <KNOWN-COUNT-NOUN>

to account for such inputs as

A CARRIER IS A SHIP.

When one of these concept-defining patterns is recognized, an acquisition procedure associated with the pattern is invoked. This procedure generally adds new facts to the system's set of wffs and generates new entries in its lexicon. The various acquisition procedures also have provisions for responding to the tutor. Response generation is accomplished through the use of preprogrammed phrases and templates.

D. Using Dialog to Aid Acquisition and Assimilation

By and large, it is unreasonable to expect tutors to volunteer all the information NANOKLAUS needs to assimilate new concepts. In particular, tutors cannot be expected to know what conclusions NANOKLAUS will draw about a newly taught concept from its previous knowledge, since they know neither the details of its knowledge nor the even more esoteric details of its assimilation procedures. NANOKLAUS must ask the tutor for the information it needs. But wording the system's questions in an understandable form involves a number of human engineering difficulties.

The acquisition of new verbs offers a cogent illustration of the problem. English verbs are highly idiosyncratic; making proper entries for them in a lexicon can be a formidable task. Among other considerations, one must ascertain a verb's passivity and transitivity, determine whether its indirect object can become the object of a FOR or TO prepositional phrase, whether it is reflexive or nonreflexive, and how the syntactic cases of its arguments may be "normalized" when the verb appears in different syntactic constructions. NANOKLAUS's tutors, of course, cannot be expected to describe verbs in linguistic terms, so the system asks a series of questions tutors can understand to elicit the same information. The following is a typical verb-acquisition exchange:

92>An officer can assign a unit
Give the -ed and -en forms of ASSIGN (for example, for GO: WENT GONE):
>>assigned assigned
If an OFFICER can ASSIGN an UNIT something, type the most general description of that something, otherwise <return>
>> (user typed a carriage return)
If an OFFICER can ASSIGN something an UNIT, type the most general description of that something, otherwise <return>
>>a task group
Can an OFFICER ASSIGN an UNIT TO a TASK GROUP ? >>yes
Could an UNIT be ASSIGNED a TASK GROUP by an OFFICER ? >>yes
Can an OFFICER ASSIGN an UNIT with a TASK GROUP ? >>no
Can an OFFICER ASSIGN a TASK GROUP ? >>no
Can an OFFICER ASSIGN ? >>no

Verbs that produce rather different answers to the above questions include GIVE, BUY, RUN (run a machine), SERVE and SUPPLY.*

E. The Knowledge Representation System

NANOKLAUS's knowledge representation system uses a many-sorted first-order logic, that combines features from [9] and [5]. The backbone of the system is a treelike data structure reflecting the

----

* NANOKLAUS does not consider modal verbs (e.g., "want" and "know"), or verbs with particles (e.g., "pick up", "preside over") or sentential objects (e.g., "The captain requested that the ship change course").

hierarchy of sorts (or classes) used by the system. This tree is "tangled" in that any given sort may have multiple ancestors. It is no accident that a sort hierarchy should serve as the primary data structure for an English-speaking acquisition system. "ISA" hierarchies are used by many natural-language processing systems; it appears that something very similar to a sort hierarchy plays a central role in the way humans organize their knowledge [7].

A fundamental task of the deductive system is to determine whether or not a given entity belongs to a particular sort. To aid in this task, NANOKLAUS maintains information about the direct descendants of each sort, such as whether some or all of the descendants of an ancestor exhaust (span) it, and whether two or more sibling sorts are mutually exclusive.

The utilization of sort information is clearly illustrated in the following dialog with NANOKLAUS:

80>A ship is a unit

You're saying that anything that is a SHIP is also an UNIT. What is an UNIT? >>a thing

You're saying that anything that is an UNIT is also a THING. Which of the following, if any, are proper subclasses of UNIT: PHYSICAL OBJECT, LEGAL PERSON, or MEASURE? >>none of them

Which of the following classes, if any, could never have any members in common with UNIT: PHYSICAL OBJECT, LEGAL PERSON, or MEASURE ? >>none

OK, now I have assimilated UNIT.

OK, now I have assimilated SHIP.

81>A ship is a physical object

You're saying that anything that is a SHIP is also a PHYSICAL OBJECT. Is ANIMATE OBJECT a proper subclass of SHIP? >>no

IS ANIMATE OBJECT necessarily composed of entirely different members from SHIP? >>yes

Do ANIMATE OBJECT and SHIP span the set of all PHYSICAL OBJECTS >>no

OK, now I have assimilated SHIP.

82>A submarine is a ship

You're saying that anything that is a SUBMARINE is also a SHIP. OK, now I have assimilated SUBMARINE.

83>A destroyer is a ship

You're saying that anything that is a DESTROYER is also a SHIP. Is SUBMARINE a proper subclass of DESTROYER? >>no

Is SUBMARINE necessarily composed of entirely different members from DESTROYER? >>yes

Do SUBMARINE and DESTROYER span the set of all SHIPS? >>no

OK, now I have assimilated DESTROYER.

84>A carrier is a ship

You're saying that anything that is a CARRIER is also a SHIP. Which of the following, if any, are proper subclasses of CARRIER: SUBMARINE or DESTROYER? >>neither of them

Which of the following classes, if any, could never have any members in common with CARRIER: SUBMARINE or DESTROYER ? >>both of them

What combination of the following subclasses, if any, together with CARRIER, span the class of all SHIPS (with or without overlapping): SUBMARINE or DESTROYER ? >>neither of them

OK, I have assimilated CARRIER.

85>A Kittyhawk is a carrier

You're saying that anything that is a KITTYHAWK is also a CARRIER. OK, now I have assimilated KITTYHAWK.

.
.
.

98>A carrier has a doctor who is a person

So CARRIERS have PEOPLE for DOCTORS.

Must a CARRIER have exactly one DOCTOR? >>no

99>The Constellation is a Kittyhawk

You're saying that CONSTELLATION is one instance of a KITTYHAWK.

OK, now I know about CONSTELLATION.

100>The JFK is a Kittyhawk

You're saying that JFK is one instance of a KITTYHAWK. OK, now I know about JFK.

1>The JFK is commanded by Moffet

ASSERTING THAT (MOFFET COMMAND JFK)

2>The length of the JFK is 1072 feet

V FUTURE PROSPECTS

At this time NANOKLAUS can be best described as a fragile proof-of-concept system still in its early developmental stage. During this coming year, we plan to greatly expand its linguistic coverage by replacing our current pragmatic grammar with Robinson's [10] DIAGRAM grammar. Once this has been accomplished and NANOKLAUS's verb acquisition package extended to accept particles and prepositional phrases, we believe NANOKLAUS can serve as a useful tool for aiding AI researchers in the construction of knowledge bases for other AI systems--a task that currently consumes an inordinate proportion of research effort.

As suggested in the introduction, one of our long-term objectives is the extension of KLAUS to knowing about diverse types of external software packages. Given knowledge of such packages, a KLAUS could serve as an agent that interacts with them on a user's behalf. To explore these possibilities, we plan in the near future to provide NANOKLAUS with the capability of using a conventional data base management system. In this configuration, a user should be able to tell NANOKLAUS about a new domain, about a data base containing information pertaining to that domain, and about the interrelationship of the two. The new system would then be able to use the data base in answering questions regarding the domain.

Our work in the area of knowledge acquisition per se has really just begun. As development proceeds, we plan to turn our attention to making provisions for learning by analogy, for acquiring and reasoning about the internal structures of processes, for dealing with causality, and for dealing with mass terms.

ACKNOWLEDGMENTS

The deduction system supporting NANOKLAUS was developed in large part by Mabry Tyson, with Robert Moore, Nils Nilsson and Richard Waldinger acting as advisors. Beth Levin made major contributions to NANOKLAUS's verb-acquisition algorithm. Paul Asente assisted in the testing of the demonstration system. Barbara Grosz, Earl Sacerdoti, and Daniel Sagalowicz provided very useful criticisms of early drafts of this paper.

REFERENCES

1. R. Davis, "Interactive Transfer of Expertise: Acquisition of New Inference Rules," Proc. 5th International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts, pp. 321-328 (August 1977).

2. T. G. Dietterich and R. S. Michalski, "Learning and Generalization of Characteristic Descriptions: Evaluation Criteria and Comparative Review of Selected Methods," Proc. 6th International Joint Conference on Artificial Intelligence, Tokyo, Japan, pp. 223-231 (August 1979).

3. J. A. Fodor, The Language Of Thought, pp. 124-156, (Thomas Y. Crowell Company, New York, New York 1975).

4. G. G. Hendrix, "The LIFER Manual: A Guide to Building Practical Natural Language Interfaces," Technical Note 138, Artificial Intelligence Center, Stanford Research Institute, Menlo Park, California (February 1977).

5. G. G. Hendrix, "Encoding Knowledge in Partitioned Networks," in Associative Networks - The Representation and Use of Knowledge in Computers, N. V. Findler, ed. (Academic Press, New York, New York 1979).

6. G. G. Hendrix, E. D. Sacerdoti, D. Sagalowicz and J. Slocum, "Developing a Natural Language Interface to Complex Data," ACM Transactions on Database Systems, Vol. 3, No. 2 (June 1978).

7. P. H . Lindsay and D. A. Norman, Human Information Processing. (Academic Press, New York, New York, 1972.)

8. T. M. Mitchell, "Version Spaces: a Candidate Elimination Approach to Rule Learning," Proc. 5th International Joint Conference on Artificial Intelligence, Cambridge, Massachusetts, pp. 305-310 (August 1977).

9. R. Moore, "Reasoning from Incomplete Knowledge in a Procedural Deduction System," MIT Artificial Intelligence Laboratory, AI-TR-347, Massechusetts Institute of Technology, Cambridge, Massechusetts (1975).

10. J. J. Robinson, "DIAGRAM: an Extendable Grammar for Natural Language Dialogue," Technical Note 205, Artificial Intelligence Center, SRI International, Menlo Park, California (February 1980).

11. P. H. Winston, "Learning Structural Descriptions from Examples," Chapter 5 in The Psychology of Computer Vision, P. H. Winston, ed. (McGraw-Hill Book Company, New York, New York (1975).