Integrating Artificial Intelligence and Graphics in a Tool for Microfossil Identification for Use in the Petroleum Industry

Peter Alan Swaby

This chapter describes an expert system for the identification of microfossils. This graphic expert system shell was designed to allow users to enter information about a fossil in pictorial form: On the basis of this information, the system selects a best-match set of fossils. By computerizing knowledge elicitation and entry into the system, it was possible to reduce development time and cost. As a result, the system is cost effective as well as powerful, flexible, and easy to use.

The system described here, called Vides (visual identification expert system), was developed to support the process of identifying microfossils, including those of Phylum Conodonta (Higgins and Austin 1985) and Phylum Foraminifera (Haynes 1981). Microfossils are the remains of small animals that lived hundreds of millions of years ago and are found in the rock layers during the drilling process of oil exploration.

Their identification serves as a guide to the geological history of the area and helps geologists estimate the likelihood that oil will be found. There are several thousand species within each phylum; distinguishing between these species is sometimes difficult, even for trained paleontologists. Oil companies can spend a million dollars a day to keep rigs operating while they wait for word from the experts; they are eager to speed microfossil identification. Traditionally, paleontologists have made use of *taxonomic keys*, which basically consist of a series of structured rules, to guide the identification process. These keys do, however, have several serious drawbacks:

First, taxonomic keys are difficult to use. They guide the identification process without incorporating all the knowledge that the user might need. In addition to the keys, paleontologists often have to consult reference books, catalogs, card indexes, and so on, during the identification process. Thus, the user frequently has to simultaneously deal with several sources of information.

Second, taxonomic keys are sequential. The nature of these keys enforces a strictly sequential identification procedure and, therefore, compels the user to choose an option—even in cases of uncertainty. Because fossils are often broken, it might not be possible to answer a question posed by the taxonomic key at an early stage of the identification procedure. To proceed with the identification procedure, the paleontologist nevertheless has to pick one of the options presented in the key, which inevitably increases the error rate.

Third, taxonomic keys are text based. The analysis of initial knowledge elicitation sessions with the expert paleontologist revealed that the process of microfossil identification is inherently a visual one: The expert would always draw a picture of fossils and their attributes to distinguish between them and only then assign verbal descriptions. Taxonomic keys do not support the visual approach a paleontologist would take given a choice.

Previous expert systems for fossil identification have been purely text based (Wiley 1987) or have only made limited use of the graphic nature of the problem (Conrad and Beightol 1988; Brough and Alexander 1986; Riedel, 1989). Although it is not difficult to translate taxonomic keys into a rule-based representation or shell, such an approach would not eliminate the inherent problems as previously described. Clearly, a system for microfossil identification should incorporate all the necessary knowledge; allow a flexible identification procedure; and let the user describe the fossils in pictorial, rather than textual, form. By combining AI techniques and graphics in an innovative way, the expert system provides such a tool.

Hardware and Software

The system was developed and currently runs on a Sun 3/260. This hardware allows identifications to be made within an acceptable time, and the high-resolution screen allows for displaying high-quality images. In addition, the availability of relevant software for Sun workstations and Unix-based computers in general proved particularly useful for the development of the image metafiles. The multiprocessing environment allows computer- intensive operations, such as image processing, to be carried out at the same time as other operations, such as text editing.

The system was developed using Lisp and IntelliCorp's knowledge engineering environment (KEE). Lisp (Steele 1984; Charniak et al. 1987) has a proven track record for AI applications such as those reported by Shortliffe (1976) and Moses (1971). KEE, as described by Fraser (1987), allows the development of systems to a high standard in a comparatively short period of time. Together, Lisp and KEE contain a large number of high-level programming primitives. KEE, in particular, provides a large number of graphic primitives that considerably facilitated the development of the user interface.

Description of the System

This section describes how the system overcomes the problems of traditional methods by applying AI techniques and graphics to the domain of microfossil identification.

Basic Solution

Incorporation of all the relevant knowledge: Ultimately, a complete knowledge base for a phylum (several thousand species) can be estimated to contain 3,500 images, in excess of 100 attribute-value tables (table 1) and more than 10,000 lines of additional information in text form. Although size as such is no problem, eliciting information on this scale and transferring it into machine-readable form would require significant time and effort on behalf of the expert and knowledge engineer. Traditional manual methods of knowledge elicitation and knowledge entry would be time consuming and, therefore, would render the system cost ineffective (Walker 1988). To avoid these problems, both knowledge elicitation and entry into the system were computerized after the initial stage (for an overview, see Williams [1990] and for a detailed description, see Swaby [1989, 1990]).

Flexible access and user-driven interaction: Choosing an appropriate form of knowledge representation ensures flexible access to, and

CAVUSGNATHIDAE	Platform Blade	Blade Position	Prominent Denticles
Cavusgnathus	present	right only	present
Adetognathus	present	right or left	present and absent
Cloghergnathus	absent	right or left	present
Clydagnathus	absent	right only	present
Rhachistognathus	absent	right or left	absent
Taphrognathus	absent	right, left, or median	present
Patrognathus	absent	right or left	posterior only

Table 1. An Attribute-Value Table for Family Cavusgnathidae.

speedy retrieval of, relevant facts from the knowledge base. After the initial knowledge elicitation sessions, it was decided that knowledge could most appropriately be modeled in the form of a table of species, attributes, and attribute values (table 1). Combined with a logic-based inference mechanism, this form of representation allows the user to access the knowledge stored in the system in a completely flexible way.

A graphic interface for visual selection: This facility allows users to describe fossils to be identified by consulting images of fossils, their attributes, and attribute values, which are stored in an image knowledge base and presented to the user on the screen. Information can be selected and entered on the basis of pictorial representations of the fossils and their attributes. This form of interaction is most appropriate for the task domain and makes the system easy to learn and easy to use.

Implementation Details

At the beginning of the knowledge elicitation process, the expert—a paleontologist—introduced the knowledge engineer to the problem of microfossil identification. The initial session was conducted in an unstructured manner: The expert explained the classification of microfossils to the knowledge engineer with the example of two genera (Gnathodus and Scaliognathus). The knowledge engineer asked questions to determine the criteria by which microfossils are classified. The session served as a basis for the knowledge engineer's first attempt to model the expert's knowledge; this model was then refined after a review by the expert. In agreement with the expert, it was decided to proceed with knowledge elicitation in a structured form: The expert would supply attribute-value tables, a description of each attribute, and a drawing and description for each attribute value as well as a drawing and a description for each fossil. When information on this scale (that is, for an entire phylum) has to be elicited and transferred into machine-readable form, manual methods would be time consuming and expensive. Consequently, it was decided that it would be an advantage to automate this process. The images supplied by the expert are translated into machine-readable form by an image scanner. Special program routines now allow images to be automatically processed and enhanced. The attribute-value tables and additional text descriptions are supplied in electronic form (through electronic mail) by the expert. The knowledge engineer checks the information provided at the data-entry stage, using routines to ensure validity and completeness. Further routines provide feedback for the expert with regard to the discrimination power of the knowledge provided for each taxon.

The approach to knowledge representation was to store text-based information and images in separate parts of the system. Textual information for each taxon is stored in a knowledge metafile. Fossil groups are classified in a tree structure from the phylum level down to the species level. Consequently, the metafile contains all the attributes and attribute values and their descriptions and all the related taxa with their attributes and attribute values, which describe them, and other relevant information (figure 1). The knowledge in the metafile provides the basis from which frames (figure 2) are generated (by the knowledge base building tool; discussion to follow) in the course of a consultation. A set of reduced metafiles contains the knowledge of higher-level taxa. Their characteristics are more distinct, and the fossils can be identified without working through attributes and attribute values. Therefore, the user can differentiate between the fossils by simply comparing their images. The images for each taxon are stored in an image metafile. These metafiles contain the main image, all the species images, and images of their attribute values (figure 3). Again, the image metafiles for the higher-level taxa contain only images of related taxa because these can be identified by simply matching fossils to their appropriate image. The conversion of metafiles into knowledge bases is done by a Lisp program—the knowledge base building tool. This program creates all the code and data structures that the expert system needs. It also produces all graphic windows, complete with images and descriptions. Once a knowledge base has been completed, it is saved and can be used by the expert system. During consultation, the user selects images to enter information about the object to be identified. The knowledge bases presented to the user can be generated from the metafiles in two different ways. One way is to convert the information only when needed—this approach would be efficient in terms of storage space but would increase the system response time as much as several minutes in some cases. The second way is to convert the information beforehand and store it in (KEE) knowledge bases; this approach is less efficient in terms of storage space but allows the user to move quickly between information in different parts of the classification

```
genus name
        attribute name
            attribute value
                description
            .
            description
        .
        species name
            author
            brief description
            attribute
                value
                .
            .
            basin
                environment
                    range
                    .
                .
            .
            detailed catalogue description
        .
        knowledge base type
        group name
            subgoup name
```

Figure 1. Knowledge Metafile Structure.

structure. In the interests of usability, the second approach was chosen. As well as the knowledge base building tool, a column building tool, a chart building tool, and a dictionary building tool have been written to produce other important parts of the system.

The inference mechanism matches the description entered by the user with the knowledge in the knowledge base and retrieves a series of possible solutions. For the system in question, the inference mechanism had to cope with three special constraints. First, fossils are often broken. Therefore, paleontologists are not always able to identify all the features required to reach a simple conclusion. Second, in some

name	delicatus
range	Upper Typicus
carina	unexpanded
inner paraper	long
outer parapet	absent

Figure 2. Frame Structure.

main image

species image
.
.
attribute value image
.
.

Figure 3. Image Metafile Structure.

cases, it is possible to distinguish between fossils on the basis of a subset of attribute values. Therefore, it should not be necessary to enter an attribute value for each attribute to uniquely describe a fossil. Finally, a paleontologist is not always able to precisely identify attribute values. Therefore, the system has to accept less precise information; for example, an attribute is one of a subset of values. Another important consideration was to preserve the transparency of inference: The expert and the knowledge engineer agreed that because the system was designed as a tool for paleontologists, it must be possible for them to follow or reconstruct the working of the inference mechanism. The usability of the system would be seriously threatened if the inference mechanism were opaque: Users cannot be expected to trust a system whose working they do not understand. Furthermore, the user will not be able to identify implausible solutions delivered by the system (which might occur, for instance, because of description errors during the consultation). The inference mechanism is based on Boolean logic: Attribute

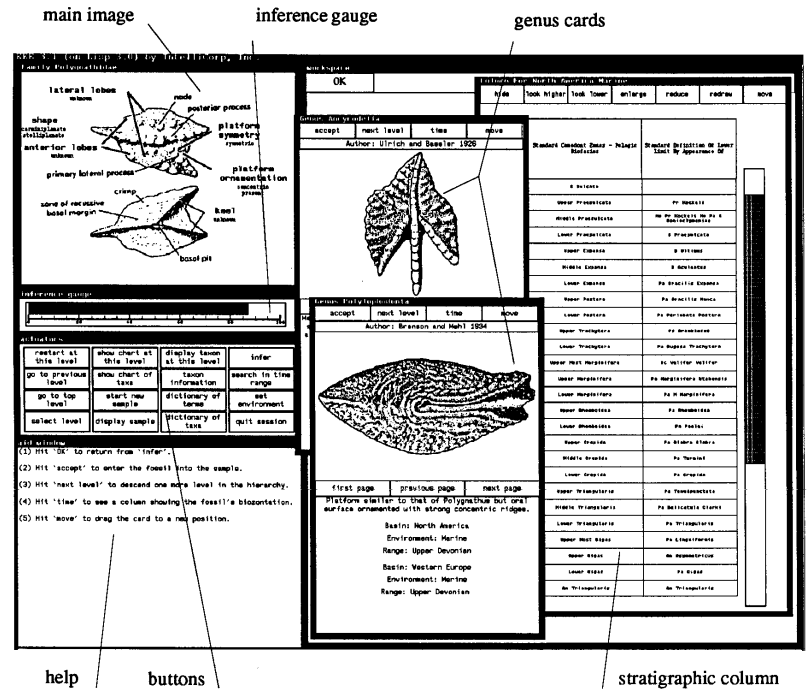

Figure 4. Screen Layout.

values entered are used to select a subset of fossils. Information is processed in a simple set-covering way to match possible fossils within the group.

An Example Session

The individual parts of the graphical interface (figures 4–9) are described here in the order the user would encounter them in the course of a consultation. Initially, the user would either start at the top level or select an appropriate level—through a cascading menu—using the select-level button in the actuator window (see figure 6). Once the level is chosen, the user is in a position to begin the identification process; help is given at each stage in the form of a list of possible options in the aid window (see figure 4). The identification is done by consulting the image window that contains a picture of the fossil shape relevant to the chosen level. All attributes of the fossil appear as labels (see figure 5). The user can choose any of the labels with "unknown" written under them, which causes a card with a description of the attribute and a series of cards with descriptions and pictures of the attribute values to

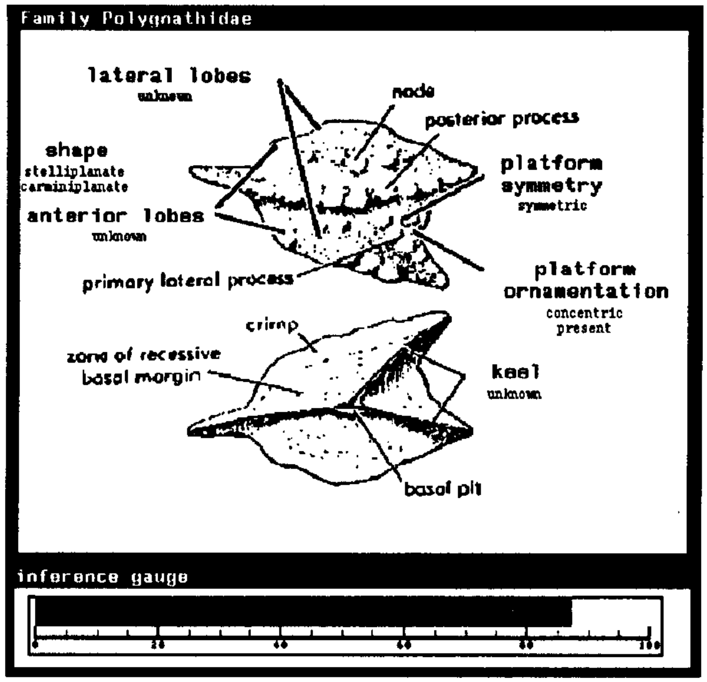

Figure 5. Main Image and Inference Gauge.

be displayed in the work-space window (figures 7 and 8). These descriptions might contain technical terms that the user might be unfamiliar with. However, at any point, a dictionary of terms and pictures can be consulted by selecting the dictionary button in the actuator window.

A fossil description is built by selecting the most appropriate cards in the work-space window and clicking on their accept buttons (see figures 7 and 8). The chosen values are written below their labels in the image window, and the inference gauge is updated (figure 5). The inference gauge gives a measure of the current degree of discrimination as a percentage, shown graphically as a proportional-length shaded bar: Zero percent means that the system cannot select any species as more likely to fit the information given at this point than another; 100 percent means that the system is certain that the information matches only one species. When the user is satisfied that the description entered accurately describes the fossil to be identified, the infer button is pressed. This action will cause the system to display the fossils that match the information supplied (see figure 9). The time range for a

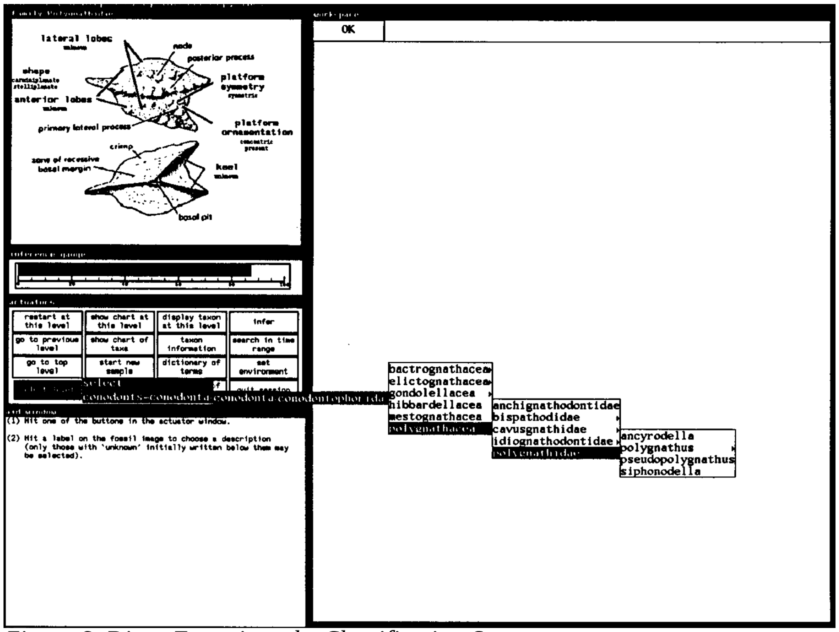

Figure 6. Direct Entry into the Classification Structure.

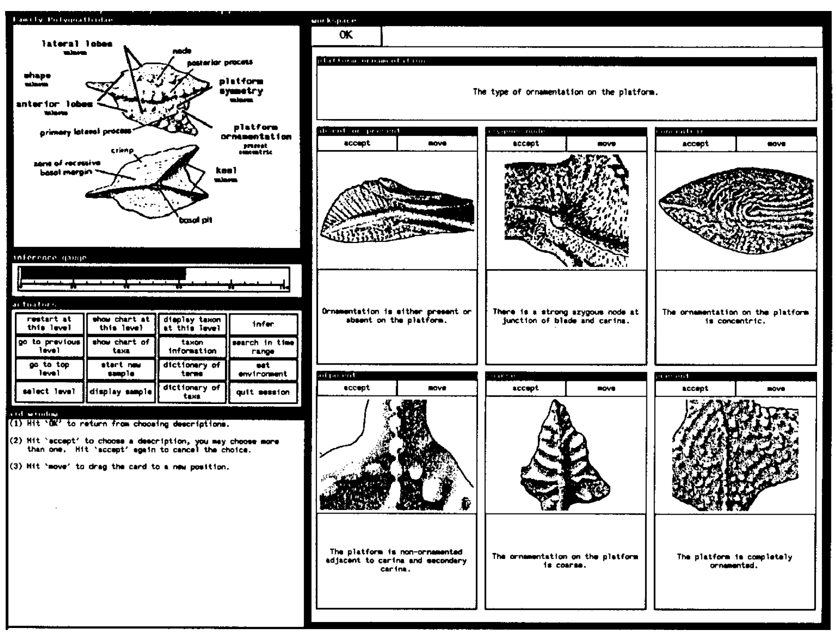

Figure 7. Description of Platform Ornamentation.

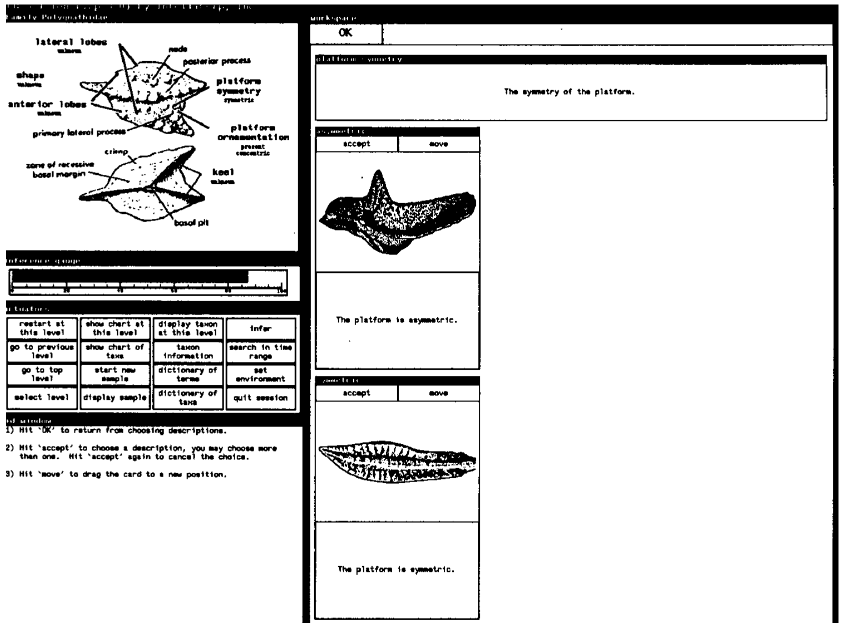

Figure 8. Description of Platform Symmetry.

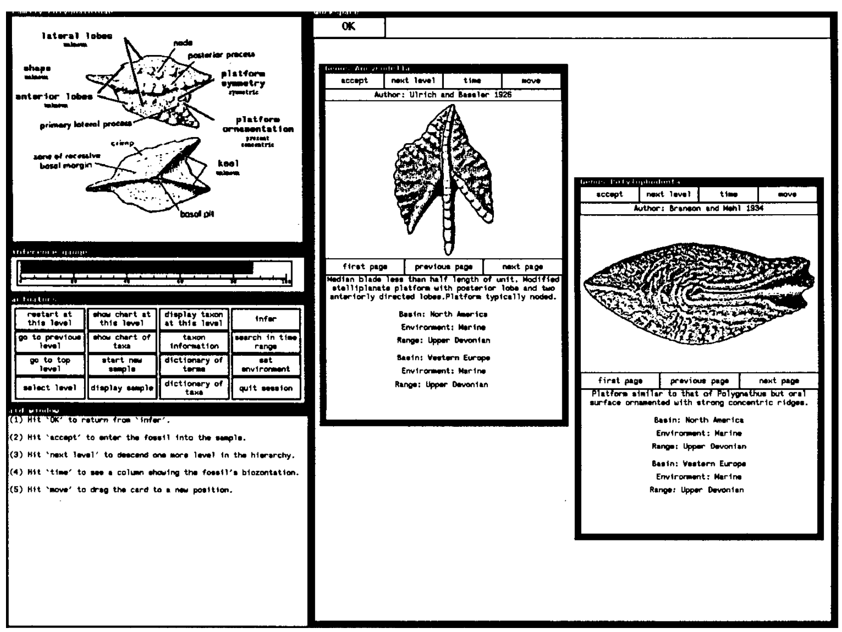

Figure 9. Inference of a Solution.

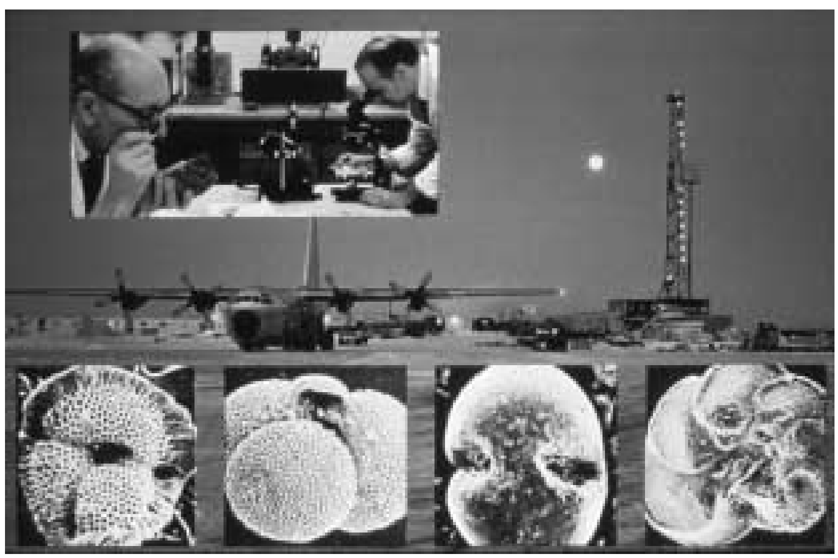

fossil can be shown on a stratigraphic column by pressing the time button on the fossil's card (see figure 4). At this point, the user can record a fossil into the current sample set by pressing the accept button on the fossil's card or proceed to a lower level in the hierarchy by choosing the next-level button.

Evaluation

The system's performance has constantly been monitored by the expert paleontologist. Furthermore, a detailed study with three potential users working—after a few hours training—through a set of benchmark tests took place. According to the expert, the level of accuracy exhibited by the subjects on the identification tasks would take several months of training to achieve if traditional methods were used. The users filled in an evaluation checklist (Ravden and Johnson 1989) after completing the benchmark tests. Analysis of the checklist and an informal discussion of the results with the users revealed that even though the overall usability of the system was favorably rated (especially the ease of input and the transparency of the inference mechanism), some features of the interface could be improved on. The features mentioned were a color display (instead of monochrome), better integration of the dictionary tool, more detailed labeling of some windows, and the provision of a reverse action–undo facility.

Discussion

This system was developed using techniques from the fields of AI and graphic interface design to provide a more advanced tool for paleontologists than has been attempted to date. By making use of the graphic nature of the problem of fossil identification, the system takes a novel approach. Information is entered at each level of interaction—through a mouse on the graphic display—by choosing the most appropriate pictures and descriptions that are presented. This approach makes the system easier to use than others that have relied on textual descriptions and have required keyboard input (Wiley 1987; Conrad and Beightol 1988; Brough and Alexander 1986; Riedel 1989). Furthermore, the system differs from previous ones in that all the knowledge is being elicited and entered for complete phyla and not just a few test cases to provide a demonstration.

Eliciting and representing knowledge on the scale necessary for this system meant that traditional approaches had to be carefully reviewed and automated wherever possible. The cost effectiveness of expert systems in geological exploration has been questioned recently (Walker

1988); by computerizing several steps of the traditional process, it was possible to build this expert system from scratch and develop a knowledge base for Phylum Conodonta fossils in just 14 person-months (12 months for the system builder and 2 months for the expert). Furthermore, a knowledge base for Phylum Foraminifera fossils down to the genus level was produced in only 4 person-months (2 months for the system builder and 2 months for the expert). The cost of oil exploration, in which the task of microfossil identification plays a central role, is huge, typically several million dollars to drill each well.

To illustrate the cost savings associated with computerization, British Petroleum was looking for oil in Papua, New Guinea, and had to halt the drilling process at one rig while the fossils they found were identified. This process involved flying out British Petroleum's expert for this fossil group to the rig. When the expert arrived, he found that a mistake had been made, and the fossils belonged to a group with which he was not familiar; another expert had to be flown to the rig. The delay cost the company $15 million. Had an expert system complete with all microfossil knowledge bases been at the rig, a less specialized paleontologist could have made the identifications immediately.

The knowledge required to identify fossils is held by only a few experts, and fossils that are difficult to identify have to be flown to them at the research center. Through the expert system, this knowledge is captured and can be distributed to less specialized paleontologists at the decision site, for example, the oil rig. It would take a non-expert paleontologist three to four months of training to identify Phylum Conodonta fossils using traditional methods, but the evaluation showed that they can correctly identify fossils in approximately 40 minutes after a few hours of training on the expert system. Therefore, it would be feasible for paleontologists at the drilling site to use the system after a short training period. Fast, accurate identification at the drilling site saves time and considerable cost in petroleum exploration. Compared to the potential savings that could be made during the drilling operation, the investment in the development of the expert system is insignificant. Furthermore, the system helps preserve the scarce and valuable resource of the experts' knowledge.

Conclusions and Further Work

This system can still be further developed and improved, and several major areas of work have not even been considered. The system could be improved if, once complete, it was recoded using different software, such as C and X-windows: It would run faster, require less memory and disk space, be more portable, and cost less in terms of software licences. These needs would be prerequisites for it to be deployed at the rig site. A knowledge elicitation tool might be developed that would allow the knowledge engineer or even the expert to interactively enter the knowledge in the knowledge base. Such a tool could check for logic errors and knowledge base consistency and provide immediate feedback, reducing the time needed to enter the knowledge. Finally, it might be possible to extend the system to automatically identify fossils—using input directly from a microscope. Even if this could only be done for the simpler fossil shapes, leaving the more complex ones to be identified by paleontologists, it could potentially save a lot of time, increase speed, and further decrease the cost of the identification process.

Acknowledgments

The author would like to thank Richard Howarth of British Petroleum for his support and encouragement, Alan Higgins and Fred Banner of British Petroleum for their expertise and help, Angela Sasse of the University of Birmingham for her assistance in writing this chapter, and Ib Sorensen of British Petroleum for suggesting many improvements to the chapter.

References

Brough, D. R., and Alexander, I. F. 1986. The Fossil Expert System. *Expert Systems* 3(2): 76–83.

Charniak, E.; Riesbeck, C.; McDermott, D. V.; and Meehan J. R. 1987. *Artificial Intelligence Programming.* Hillsdale, N.J.: Lawrence Erlbaum.

Conrad, M. A., and Beightol D. S. 1988. Expert Systems Identify Fossils and Manage Large Palaeontological Databases. *Geobyte* 3(1): 42-46.

Fraser, J. 1987. *Some Aspects of Programming in KEE.* Edinburgh, Scotland: Artificial Intelligence Applications Institute, University of Edinburgh.

Haynes, J. R. 1981. *Foraminifera.* London: Macmillan.

Higgins, A. C., and Austin, R. L. 1985. A Stratigraphical Index of Conodonts. Chichester, England: Ellis Horwood.

Moses, J. 1971. A Macsyma Primer, Mathlab Memo, 2, Computer Science Lab., Massachusetts Institute of Technology.

Ravden, S., and Johnson, G. 1989. *Evaluating Usability of Human-Computer Interfaces.* Chichester, England: Ellis Horwood.

Riedel W. R. 1989. Identify: A Prolog Program to Help Identify Fossils. *Computers and Geosciences* 15(5): 809–823.

Shortliffe, E. H. 1976. *Computer-Based Medical Consultations: Mycin.* New

York: American Elsevier.

Steele, G. L., Jr. 1984. Common Lisp: The Language. Pensauken, N.J.: Digital.

Swaby, P. A. 1990. Towards Large Knowledge Bases: Automating Knowledge Elicitation, Knowledge Verification and Knowledge Base Building. In *Research and Development in Expert Systems VII, Proceedings of Expert Systems 90,* 215–225. Cambridge: Cambridge University Press.

Swaby, P. A. 1989. A Graphical Expert System for Microfossil Identification. In *Research and Development in Expert Systems VI, Proceedings of Expert Systems 89,* 240–252. Cambridge: Cambridge University Press.

Walker, M. G. 1988. Expert Systems in Geological Exploration: Can They Be Cost Effective? *Geobyte* 3(3): 18–23.

Wiley, P. A. 1987. To Evaluate the Use of the Expert System Builder Crystal for use in Conodont Identification, Master's thesis, Kingston Polytechnic.

Williams, S. 1990. Identifying Fossils by Computer. *Science* 248:1080.