CHARLES FENLY*

Assistant Chief, Descriptive Cataloging Division
Library of Congress
Washington, DC

Technical Services Processes as Models for Assessing Expert System Suitability and Benefits

ABSTRACT

The rudiments of a strategy for assessing the appropriateness of using an expert system in a given domain are presented. The assessment process involves comparing the characteristics of the domain with various suitability criteria and identifying potential benefits from an expert system in the domain. Two library technical service functions, descriptive cataloging and shelflisting, are used as models illustrating this assessment process. Based on organizational factors specific to the Library of Congress, series work (a subset of descriptive cataloging) and shelflisting appear to be suitable and beneficial candidates for expert system development efforts.

INTRODUCTION

Two questions that might be posed about the introduction of expert systems technology into the technical service workplace are: *Can* expert systems be applied to library technical service processes? and, if so, *Should* expert systems be applied to library technical service processes? Of these, the second is by far the more interesting and challenging question.

*The views expressed are those of the author and do not reflect the official policies of the Library of Congress.

The first question may be answered fairly easily. An examination of the literature of expert systems suggests rather clearly that, given the will, time, and resources, expert systems could be applied in some way to such activities as cataloging, classification, acquisitions work, serials management, and the like.

But even if it is accepted that expert systems could be implemented in technical services operations, the question remains, should they be? There is no single, conclusive answer to this question. However, there are strategies for approaching this question that an organization can use to help make rational decisions about whether an expert system has a place in its own operations. This paper will present some of the fundamentals of such a strategy.

DECIDING ON AN EXPERT SYSTEM

As a technical services manager, this author believes strongly that decisions related to implementation of expert systems technology in the technical services (or any other) workplace should be based on sound management decision making. This may sound obvious, and perhaps it would be obvious if the topic of discussion were something other than an aspect of artificial intelligence (AI) technology. But AI technology carries with it such a degree of fascination and, one might even say, glamour, that the possibility of trying to introduce such technology for its own sake rather than for sound management reasons is present with respect to expert systems in ways that might not be the case with other technologies.

What is it that prompts a consideration of using an expert system in technical processing? First, there is an increasingly wide awareness among technical services librarians that computer programs have been developed which exhibit human-like reasoning, which may be able to learn from their mistakes, and which quickly and cleverly perform tasks normally done by scarce and expensive human experts. Further, it is widely recognized that automation has paid off in a big way in technical processing operations in the past: through creative use of computing, marvels of information storage and retrieval and resource sharing have been achieved. It is therefore natural that technical services librarians would wish to assess whether this newer technology has the potential to confer similar benefits.

Upon further investigation of the matter in the literature, an increased understanding of the realities and limitations of expert systems technology and a greater awareness of what expert systems can and cannot do might lead to some decline in enthusiasm. While it may be true that expert systems have been programmed to solve complex problems, this has tended to require considerable expenditure of time and money, and even after such expenditure, many systems have never gone into production because of reliability problems. And though research is underway to improve the processes by which systems learn, at present, the acquisition of knowledge by an expert system is one of the most difficult aspects of system development—a major bottleneck rather than one of the strengths of the technology (Rolston, 1988, pp. 157-67). Nevertheless, there is something rather compelling about this technology, so that even while recognizing that it is not a panacea, we may continue to have a strong interest in examining more closely whether there are prospects for using expert systems in our organizations.

According to Greene (1990, pp. 48-59), artificial intelligence technology has become so well integrated in some Japanese organizations that it can function as a frequently used tool of problem solving and work improvement. Clearly, this is not the situation that prevails in libraries today. Artificial intelligence technology is relatively unknown and may even be viewed as somewhat exotic, and an effective way to get to know it better is to investigate its potential usefulness.

A much less defensible approach to learning about expert systems, however, is to embark hastily upon an expert systems development project based on the premise: "Let's think of something that we can develop an expert system to do." The objection to this approach is a practical one: It is too likely to result in projects that go nowhere, in systems that do not produce useful results, that is, "toy systems."

This does not imply that there is anything wrong with developing a small expert system. There may be definite benefits to be gained from implementing a small system which deals effectively with a real problem which happens to be small in size. As our familiarity with the technology increases, and as more powerful and user-friendly development tools become available, it may become increasingly common for domain experts to engage in their own knowledge engineering to develop such systems to help them do their work. But this appropriate use of expert systems technology is quite different from projects whose end result is a "demonstration prototype": a small system which is small because (a) it deals with a tiny piece of a large domain, with no clear plans to expand it to the point where it can address a meaningful subset of the domain; or (b) it is a shallow and superficial cut at a deep and complex problem. Projects such as these do not confer upon an organization the kind of benefits which expert systems have the potential to yield.

There are no doubt many different strategies which could be proposed for assessing the appropriateness of implementing an expert system in a given domain in a particular organization. The discussion which follows will suggest one possible approach. Though the approach described has general applicability, the discussion will relate the issues addressed to specific technical processing activities.

Assessing the Appropriateness of an Expert System

In 1987, Howard Harris and I conducted an investigation of the feasibility of applying expert systems to technical processing operations in Processing Services (now known as Collections Services) of the Library of Congress (Fenly & Harris, 1988). Since neither of us was an expert system expert, we began with an extensive literature review to gain a better understanding of the technology. That survey of the literature convinced us that expert systems potentially have great power. But it also convinced us that genuine expert systems, with the depth and power to solve substantial and meaningful problems, are time consuming and costly to develop and that expert system development projects have uncertainties associated with them that would probably not be tolerated in connection with traditional data processing initiatives. Thus, it was clear that there have to be other reasons for implementing an expert system besides the fact that it would be intellectually stimulating to do so.

And, in fact, there are other reasons: namely, the potential benefits to be derived from a successful implementation of an appropriate expert system. Some of the potential benefits of expert systems that have been described in the literature are these:

- expert systems can make scarce expertise more widely available within the organization, thereby helping nonexperts achieve expert-like results;
- they can free human experts for other activities besides repeatedly solving the problems which an expert system could address;
- they can promote standardization and consistency in the solving of relatively unstructured tasks;
- they can enhance organizational effectiveness and efficiency by making readily available solutions to difficult problems which might otherwise require time-consuming research or consultation with experts to solve;
- they can provide a means for capturing and storing valuable knowledge that might be lost if an employee with scarce expertise left the organization;
- they can provide a means for long-term retention of complex knowledge, since machine knowledge does not deteriorate with time or disuse in the same way that human knowledge does;
- they can perform, at a consistently high level, tasks which humans might perform inconsistently due to fatigue or loss of concentration (Beerel, 1987, pp. 84-85; Olsen, 1989, pp. 121-22; Waterman, 1986, pp. 12-13).

As stated above, these are potential benefits to be derived from successful implementation of an *appropriate* expert system. That raises another question: What are the criteria for assessing whether a particular domain is *suitable* for an expert system?

Assessing the Suitability of an Expert System

The following list of suitability criteria (a slightly modified version of the list used in the 1987 Library of Congress study) was based on work by Prerau (1985). It must be emphasized that the list given here is only representative, not exhaustive. It is intended to give a flavor of the characteristics of an expert system type of problem. For a detailed and comprehensive discussion of this important topic, see Prerau (1990).

Note that these are referred to as "suitability" criteria, not "feasibility" criteria. The fact that it might be feasible to apply AI programming techniques to a problem does not in itself make that problem a suitable domain for an expert system. Conventional data processing techniques or nonautomated tools such as manuals, flowcharts, or decision logic tables may be more appropriate ways of addressing a given problem or task.

Selected Domain Suitability Criteria

1. Tasks to be performed and problems to be solved in the domain require expert knowledge, judgment, and experience. In other words, problems in the domain are nontrivial, and experienced people perform the work at a significantly higher level than novices.
2. Tasks and problems in the domain require primarily symbolic (rather than algorithmic) reasoning and require the use of heuristics. Otherwise, more familiar and possibly more efficient conventional data processing techniques might be more appropriate.
3. Tasks and problems in the domain have appropriate depth. In practice, tasks to be performed might typically take an expert a few minutes to a few hours to perform. A domain which lacks depth is not a good expert system domain for at least two reasons: First, if tasks typically take only seconds to perform, the work might actually be slowed down by the time required to interact with an expert system; second, users are likely to become bored with and discontinue using a system that answers only simple questions. On the other hand, if a domain is so deep that tasks take many hours to complete, an expert system in the domain might be unmanageably large and unacceptably slow and expensive.
4. The domain is relatively narrow, well bounded, and self-contained. Since an expert system should deal with problems of meaningful depth, a domain which is extremely broad or unbounded could

overwhelm a development effort. This potential problem can be mitigated, however, if a large domain can be segmented into manageable parts.

5. Some degree of incorrect or nonoptimal results can be tolerated. This is important because expert systems are subject to producing unreliable or invalid results due to certain inherent limitations of the technology (Waterman, 1986, p. 29; Hollnagel, 1989, pp. 33-35). Furthermore, expert system knowledge engineering is subject to the law of diminishing returns to scale. That is, the incremental utility of the system increases by ever smaller amounts, as additional costs are incurred to improve system performance by increasing the percentage of domain knowledge embodied in the system. Eventually, a point will be reached where the marginal cost of adding further knowledge will exceed the resulting marginal increase in the usefulness of the system, at which point knowledge acquisition should cease. The system will contain less than complete domain knowledge at this point (Kang & Levy, 1989, pp. 242-43).
6. The domain is fairly stable, with the need for the task projected to continue for several years, with changes tending to be gradual and evolutionary, and with no radical changes which would redefine the task being planned. This is important because of the anticipated amount of time associated with development of a substantial system. It might be very hard to evaluate during the development process the performance of a system in a highly volatile domain. Needless to say, it is important that the system still be useful and relevant at the time it is ready to go into production.
7. There are recognized experts working in the domain who would be willing and available to participate in a development project. These experts would normally be the principal source of the expertise which is to be embodied in the system.

As already stated, the above list is not exhaustive, but it does provide some essential considerations in determining whether a task is an expert systems kind of task or whether it might better be dealt with through other approaches, such as manual processes or conventional algorithmic automated data processing.

Organizational Factors to be Considered

The criteria just discussed focus on the technical suitability of a domain as a potential expert systems candidate. In any given organization, there will also be organizational factors which would have to be considered before deciding whether to embark upon a development project. These might include such considerations as the following:

1. Is management supportive of a development project? Will management fund the project or support the seeking of funds from other sources? Can the organization's hardware, software, and professional time be devoted to a development effort?
2. Are there political objections to development of a system in the domain under consideration? Will management and staff working in that domain feel threatened or intimidated by the expert system and resist it?
3. Is there organizational support for maintaining a system once implemented? Even a seemingly stable domain might turn out, upon closer inspection, to be more volatile than one might have thought. An unmaintained system might soon begin giving wrong answers and could then be expected to fall into disuse.

It is obvious that the "wrong" answer to one or more of these questions could make a development project untenable.

DESCRIPTIVE CATALOGING AND SHELFLISTING

Now that some of the potential benefits which an organization might derive from expert systems and some criteria for determining the characteristics of an appropriate candidate for an expert system have been presented, two traditional technical services functions will be used as models of how a process might be examined against the suitability criteria and benefits. These two functions are descriptive cataloging and shelflisting. They were chosen as the models for this discussion largely because they are presumed to be more generally familiar than many other technical services functions, such as acquisitions and serials management, which may be performed rather differently at different organizations. It should nevertheless be emphasized that decisions about expert system development have to take the environment of the specific organization where they are intended to be used into account, and that will be reflected in what follows. In the present discussion, the organizational realities are those of the Library of Congress.

Descriptive Cataloging as an Expert System Domain

Descriptive cataloging is the subset of cataloging activity which involves (1) providing a bibliographic description of an item sufficient to identify the item and to provide to a prospective user certain information necessary to make judgments about its usefulness, and (2) formulating uniform access points to enable the potential user to retrieve the bibliographic record.

An examination of the library and information science literature reveals a definite interest in considering the application of expert systems technology to the rules and procedures of descriptive cataloging. This is hardly surprising, since the most common expert system knowledge base building block is the rule, and descriptive cataloging is certainly rule oriented; indeed, with its basis in the *Anglo-American Cataloguing Rules,* second edition *(AACR2)* (Gorman & Winkler, 1988), it is one of the most codified domains in librarianship. A particular focus of interest for purposes of suggesting hypothetical expert systems or building prototype systems has been chapter 21 of *AACR2,* which deals with the process of choosing access points. Upon cursory examination, this chapter appears to lend itself to the formulation of many rules in the form exemplified by the following:

> IF court rules govern a single court
> THEN main entry is the heading for the court (Rule 21.34A, Modified from *AACR2,* 1988, p. 364)

Attempts to develop a knowledge base built in such a fashion have tended, however, to produce unconvincing results. An example of an *AACR2*-based system which has been described in the literature is CATALYST (Gibb & Sharif, 1988). This system was developed using the PC-based expert system shell, ESP-Advisor. A feature of this shell, called "text animation," facilitates the conversion of existing documentation into an expert system knowledge base. CATALYST works by presenting the user with various menus; the user is expected to indicate a choice, which the system then uses to consult the knowledge base and either advance to the next level in the decision tree or provide an answer to the problem being addressed. In the report on this system, several examples of such menus relating to choice of main entry heading are provided. An examination of these examples prompts questions about the probable usefulness of this system. It seems likely that the appropriate menu choice will often not be evident to the novice with limited cataloging knowledge. Though there is a help facility, the decision to consult it depends on recognizing what one does not know; this is often far from obvious when an inexperienced person is dealing with such complex matters as choosing a bibliographic access point. And the experienced cataloger, if he or she does not already know the right answer, will probably want to read the rules carefully in order to understand the correct approach in its context, as opposed to relying on the skeletal information provided by the help facility.

The problem at work here is one that several writers have pointed out: The expertise in cataloging is not explicit in the rules; rather, it is implicit in the heuristics employed by the experts who do the work (Davies, 1986, p. 72). Consulting *AACR2* is not synonymous with descriptive cataloging: "Like most professional handbooks, it is written for those who already know" (Hjerppe et al., 1985, p. 12). In fact, this problem is noted in the report on CATALYST, making slightly puzzling the authors' conclusion that "[CATALYST's] value as an assistant is yet to be assessed but it seems likely that it can contribute to both educational and operational environments" (Gibb & Sharif, 1988, p. 70). Another possible conclusion might have been that development of an expert system in descriptive cataloging which possesses genuine expertise would require very extensive knowledge engineering, and is therefore a problem of a completely different order of magnitude from that of using an expert system shell to recast the cataloging rules into an automated format.

Thus, the appropriateness of applying an expert system to a particular domain should not be assumed too hastily. Descriptive cataloging is rule based and expert systems are frequently rule based, but this apparent similarity is by no means adequate evidence that the descriptive cataloging rules constitute a suitable expert systems domain. A decision that a domain is right for an expert system is better arrived at through a careful comparison of the characteristics of that domain to suitability criteria such as those discussed above.

In the Library of Congress study (Fenly & Harris, 1988), such comparisons were made in a number of domains. The following is an example of the results of such a comparison with respect to the domain of descriptive cataloging.

1. Do the tasks to be performed and problems to be solved in this domain require expert knowledge, judgment, and experience? This question can be answered confidently in the affirmative. There are marked differences in performance between the novice and the experienced individual in this domain, and the time required to achieve performance levels characteristic of the best practitioners is likely to be measured in years. Thus, this is an expert domain.
2. Do the experts in this domain use symbolic reasoning and heuristic problem solving? Again, the answer is yes. This is particularly the case in subsets of the domain involving complex relationships or research, such as series work or work involving formulation of complex name headings or uniform titles.
3. Do the tasks to be performed possess the desired degree of depth? The answer here is not so obvious. Although the full process of completing the descriptive cataloging portion of a particular bibliographic record might fit neatly into the "few minutes to few hours" time frame, the process in practice consists of a number of discrete steps, and many of the necessary decisions are usually made by an experienced individual almost as quickly as he or she can examine the item being cataloged, and certainly in less time than would be required to interact with an expert system. Certain

subtasks of descriptive cataloging, however, are intricate enough in themselves to satisfy this criterion. A good example is series work. Since this subset of the descriptive cataloging domain will be the focus of further attention later in this paper, a brief discussion of series work and what a series expert system might do will be presented at this point.

A Series Expert System

The Fenly and Harris investigation at the Library of Congress suggested that series work is the aspect of descriptive cataloging most likely to require a disproportionate amount of consultation to resolve unusual problems. In fact, such consultation was involving so much of the attention of certain experts in the Office for Descriptive Cataloging Policy that the office embarked on a special training program to increase the number of series experts within the monographic cataloging sections. Several factors make series work uniquely challenging, including the problems of seriality, the number and complexity of series-related rules and procedures, and the difficulties that stem from the need to relate newly received items to existing series, many of which were established under different rules and practices from those now in place.

An expert system which would help address these problems would include the knowledge and heuristics which the best experts apply to deal with these troublesome matters. The system would assist the user in pinpointing the nature of the problem, perhaps through the use of increasingly detailed levels of menus. It would be capable of asking for information needed to evaluate the problem, and it would be able to recommend a solution or recognize that it lacked adequate knowledge to solve the problem. As a by-product of containing the facts and heuristics associated with series work, it would be capable of assisting in the establishment of a new series, including determination of proper form of headings, references, and treatment.

4. Is the task relatively narrow, well bounded, and self-contained? Our investigation at the Library of Congress convinced us that the domain of descriptive cataloging as a whole is much too broad for an expert system which attempts to cover the full range of tasks at an adequate level of depth to be appropriate. It is therefore important to subdivide this domain in order to focus on a narrow subset of problems so that a realistically deep expert system can be contemplated. Series work constitutes such a subdivision.
5. Can some degree of incorrect or nonoptimal results be tolerated? Traditionally, a high degree of accuracy in adherence to cataloging rules and procedures has been considered the norm. In the present environment of automated storage and retrieval of bibliographic information, accuracy and consistency are as important as ever. An

expert system that delivered wrong answers too often would thus be unacceptable. Unfortunately, as noted above, expert systems are subject to the law of diminishing returns with respect to fine-tuning their level of performance beyond a certain point. This poses a challenge to the would-be developer of an expert system in cataloging. Unless the system can be fine tuned to yield results of acceptable accuracy, it will either never be implemented or will quickly fall into disuse. It therefore becomes important to ask the question: What is an acceptable performance level? Rolston (1988, pp. 213-15) provides a useful perspective on this question. A primary purpose of an expert system is to distribute an expert's knowledge to non-expert users. Therefore, a system's effectiveness should be evaluated not by comparing its results to some theoretical model of perfection but by comparing its performance to what the intended users would achieve without the system's help. Viewed from this perspective, it is reasonable to assume that an expert system of an acceptable performance level could be developed in the domain of series work, though it may be hard to judge in advance how much effort would be required to attain that level.

6. Is the domain fairly stable? Are significant changes anticipated in the near future? These are most important questions because of the anticipated length of time required to bring into production a substantial expert system application. Some years ago, it would have seemed rather obvious that this was a reasonably stable domain. At present, however, it appears that environmental forces, chiefly economic, may have the effect of introducing increased volatility into this domain. In the face of budgetary constraints leading to reduced staff levels and growth of backlogs of uncataloged materials, serious attention is being given to descriptive cataloging simplification. This could have the effect of bringing about changes to existing practices, which could significantly complicate a system-development effort mounted in the near future.
7. Are there recognized experts working in the domain today? There are indeed recognized and articulate experts available to lend their knowledge and experience to a development effort.

The process of evaluating the domain of descriptive cataloging against the suitability criteria thus yields somewhat mixed results. A number of the criteria appear to be well satisfied, with those relating to appropriate task depth and domain breadth seeming to be satisfied best by one of the more complex subsets of the domain, such as series work. On the other hand, due to the diminishing returns problem in connection with expert systems development, it may be hard to predict in advance how much effort (and therefore cost) will be required to implement a system which will demonstrate acceptable accuracy levels, and any development effort mounted in the near future might be subject to being hampered by possible changes in practice in the domain.

On balance, if it is assumed that these two concerns can be satisfactorily addressed, it could reasonably be concluded that a complex subset of descriptive cataloging such as series work does appear to be a suitable expert system domain. In the course of the 1987 investigation, we concluded that series work was in fact one of the domains which, from among all the technical processing operations we investigated, seemed best to satisfy the suitability criteria.

If a domain seems suitable, it must then be determined whether implementing a system in that domain is likely to yield any benefits. There does appear to be the potential for benefits from a series expert system, including the following:

- As noted above, series expertise is scarce, and the system could be expected to make this scarce expertise more widely available.
- The system would free human series experts from repeatedly solving difficult series problems, thereby allowing them to turn their attention to other matters for which they are responsible.
- The amount of time-consuming research and consultation in an effort to resolve series problems should be reduced.
- Valuable knowledge and heuristics related to resolving series problems would be retained in an expert system and would continue to benefit the organization even if a human expert resigned or retired.

Since these are obviously significant benefits and since series work appears to be a suitable domain for an expert system, it would appear that this is an application worthy of serious consideration for a development effort. There is, however, one more crucial matter to consider: cost. That topic will be addressed below. First, the other major technical services function to be examined in detail in order to consider its suitability and benefits as an expert systems domain will be discussed. That function is shelflisting.

Shelflisting as an Expert System Domain

Because of the large volume of work passing through the cataloging and classification workstream, shelflisting at the Library of Congress is done by a separate section of more than sixty staff members. The principal intellectual effort of this work entails formulating a book number, known as the cutter number, which is added to the classification number provided by a subject cataloger to produce a call number unique to the item in hand. Though the cutter number is based on a simple table, in practice, the work is complicated by two factors. First, the classification schedules, which prescribe how the call number is to be structured, are extensive and complex. Not every classification number is completed according to the same formula. Second, because of the immense size of the existing shelflist, a cutter number derived from the cutter table can only be suggestive. The task of finally formulating the book number takes place at the shelflist itself, where the shelflister must fit the item now being processed into what has already been done.

Does this task constitute a suitable expert system domain? In considering that question, a conceptual model of an expert system-based approach to shelflisting is helpful. Such an approach might be based upon an expert system interacting with a database of shelflisting records. These records would contain the call number and the subset of the fields contained in a full MARC record on which the formulation of the call number depended (and, to permit fully automated shelflisting, fields for holdings information). This database might reside on a minicomputer or on CD-ROM supplemented with a dynamic database of shelflisting records formulated since the most recent issue of the CD-ROM file.

The expert system component would contain rules specifying how the cutter number should be derived in the case of each unique method of cuttering. Each rule would be linked to a database of classification numbers whose cutters are to be derived according to that rule. Thus, when the operator, in response to the system prompt, keyed in the classification number, the system would know which rule applied and could then ask for any additional data needed. The expert system could then apply its rules for actually formulating the cutter number. As part of this process, the system would consult the shelflisting record database to determine where the new record should fit, determine the correct cutter number based on that fit, formulate the shelflisting record, and add it to the database.

With this model in mind, a comparison of the domain against the list of suitability criteria can be made.

1. Is this a domain which requires expert knowledge, judgment, and experience? Because of the complicating factors already described, this is in fact a domain in which experienced practitioners perform much better than novices. A substantial program of formal training and a lengthy period of experience are required before a shelflister typically reaches a high level of proficiency in dealing with the full range of complex problems.
2. Does the task require symbolic reasoning and the use of heuristics? A superficial examination of the task would suggest that it is largely algorithmic. However, although the use of heuristic problem solving in this domain is not as great as in a domain such as series work,

the level of complexity of the work is such that it cannot be carried out by purely algorithmic procedures.

3. Does the task possess the appropriate level of depth? Because the shelflister must make decisions based upon the complicated and extensive classification schedules and upon the sometimes intricate realities of the shelflist, the task is not a trivial one which can be dispensed with in a few seconds. Thus, this criterion would seem to be satisfied.
4. Is the task relatively narrow, well bounded, and self-contained? Given the size of the classification schedules upon which the system would be dependent, it may be hard to see the domain as narrow. However, each separate shelflisting decision focuses on one small part of the schedules and of the shelflist itself. Furthermore, because of the way the classification schedules are structured, it should be possible to segment the domain for system development. In addition, though there are thousands of classification numbers, there are only a few ways to complete a call number. Thus, the domain appears to be sufficiently narrow.
5. Can some degree of incorrect or nonoptimal results be tolerated? Clearly, it is essential that call numbers be correct in the sense that the number in the cataloging record must match the number that appears on the shelved item. But, perhaps in some other respects, some nonoptimal results could be tolerated. If the number assigned to an item were slightly off the mark (for example, suppose an item by Jones in a given classification should shelve immediately after Johnson but gets put by the system immediately ahead of Johnson), this would certainly not be desirable, but a small number of such misassignments might not be excessively harmful. Furthermore, it is possible to conceive of ways to help prevent an excessive number of errors of this type. For example, two features that might be built into a system to assist in error-prevention are (a) a display of the system's results to the operator in context (for example, a display showing the newly derived shelflisting decision along with the two records that come immediately before and the two that come immediately after it); and (b) the ability to note and call to the operator's attention certain kinds of anomalies (for example, to note that although the rule it is applying calls for single cuttering, other records in that class seem to be double-cuttered).
6. Is the domain stable? It is, since changes in practice tend to be gradual and there are no significant new developments currently being planned.
7. Are there recognized experts working in the domain? There are experts with many years of experience who are articulate, capable of providing authoritative answers to the most difficult of problems, and whose expertise is widely recognized.

The comparison of the domain of shelflisting to the suitability criteria thus suggests that the domain is a potential candidate for an expert system development effort. With respect to potential benefits, there are several which might be anticipated from implementation of a properly functioning system along the lines of the model under discussion:

- If such a system could produce credible results with acceptable consistency, the exceptional labor intensity of the task as it is now constituted could be greatly mitigated. Staff could be redeployed to some of the many other pressing tasks in the organization which are not so amenable to being assisted by technology.
- As the system evolved and heuristics for dealing with some of the more unusual and complex problems were added, the number of time-consuming consultations with the most experienced experts could be lessened.
- The enormous shelflist as it currently exists has been developed over many years and embodies a good deal of implicit knowledge which may be fully understood only by a few individuals with many years of experience working in this area. If such a system could capture this knowledge, the operation would continue to benefit from the experience and expertise of these individuals even after they retired.
- Though this work does require some degree of expertise and heuristic problem solving, it is also production oriented and repetitious, so that the risk of errors and inconsistencies resulting from human fatigue is always present. An expert system would not be subject to this problem.

Thus, there appears to be the potential for truly significant benefits from a system which would function as proposed at an acceptable performance level.

Both series work and shelflisting seem to be suitable and potentially beneficial domains in which to apply expert systems technology. If it is determined that a proposed application is suitable and beneficial, and if it is assumed that organizational factors such as those noted above are not a barrier, should development work then proceed? That is certainly an option available to an organization intent on implementing an expert system. However, from a sound managerial decision-making point of view, a preferable next step would be a careful assessment of costs in relation to expected benefits.

Cost Considerations

No attempt will be made here to suggest a methodology for a cost-benefit analysis. For an organization lacking expertise in knowledge engineering, such an analysis may be difficult or even impossible to conduct "in-house." The literature of AI and expert systems offers little useful information about development costs. An additional complicating factor in a cost-benefit analysis is the intangible nature of some of the benefits sought from an expert system, such as wider dissemination of expertise and the capability to retain scarce knowledge. It may therefore be necessary to bring in a knowledge engineering consultant to assist in the analysis. This could be costly, since the consultant will presumably have a great deal to learn about the domain in order to offer sound judgments about how challenging the development effort is likely to be in order to achieve the hoped-for benefits.

Despite the difficulties, the alternative to such a cost-benefit analysis would be to proceed into a realm of considerable uncertainty. If it is true that a "small, fairly uncomplicated system" may cost $40,000 to $100,000, and that the cost of a large-scale system developed on a mainframe could exceed $1 million (Beerel, 1987, p. 61), it would seem highly advisable to undertake a development effort with the clearest possible idea in mind of what results are expected and what level of effort is likely to be required to achieve those results.

CONCLUSION

This paper has attempted to present the rudiments of a rational, businesslike strategy for identifying promising candidates for the application of expert systems technology. Two traditional and well-known library technical services functions were used as models to illustrate how such a strategy might work. Based on circumstances specific to the Library of Congress, series work (a subset of the larger domain of descriptive cataloging) and shelflisting appeared to be promising candidates for expert systems based on considerations of domain suitability and potential benefits (and pending a favorable cost-benefit analysis). No conclusions can be drawn from the foregoing about the applicability of expert systems to technical processing generally, however, since the appropriateness of implementing an expert system depends on so many organization-specific factors.

REFERENCES

Beerel, A. C. (1987). *Expert systems: Strategic implications and applications.* New York: Halsted Press.

Davies, R. (1986). Expert systems and cataloguing: New wine in old bottles? In F. Gibb (Ed.), *Expert systems in libraries* (Proceedings of a conference of the Library Association Information Technology Group and the Library Association Research Group, November 1985) (pp. 67-82). London: Taylor Graham.

Fenly, C., & Harris, H. (1988). *Expert systems: Concepts and applications.* Washington, DC: Cataloging Distribution Service, Library of Congress.

Gibb, F., & Sharif, C. (1988). CATALYST: An expert assistant for cataloguing. *Program: Automated Library and Information Systems, 22*(1), 62-71.

Gorman, M., & Winkler, P. W. (1988). *Anglo-American cataloguing rules.* (2d rev. ed.). Chicago, IL: American Library Association.

Greene, R. (1990). *Implementing Japanese AI techniques: Turning the tables for a winning strategy.* New York: McGraw-Hill.

Hjerppe, R.; Olander, B.; & Marklund, K. (1985). *Project ESSCAPE: Expert systems for simple choice of access points for entries: Applications of artificial intelligence in cataloging.* Linköping, Sweden: Linköping University, Department of Computer and Information Science and University Library.

Hollnagel, E. (1989). The reliability of expert systems: An inquiry of the background. In E. Hollnagel (Ed.), *The reliability of expert systems* (pp. 14-36). Chichester, West Sussex, England: Ellis Horwood.

Kang, K. C., & Levy, L. S. (1989). Software methodology in the harsh light of economics. *Information and Software Technology, 31*(5), 239-250.

Olsen, P. R. (1989). Safety and risks in the use of financial expert systems. In E. Hollnagel (Ed.), *The reliability of expert systems* (pp. 119-133). Chichester, West Sussex, England: Ellis Horwood.

Prerau, D. S. (1985). Selection of an appropriate domain for an expert system. *AI Magazine, 6*(2), 26-30.

Prerau, D. S. (1990). *Developing and managing expert systems: Proven techniques for business and industry.* Reading, MA: Addison-Wesley.

Rolston, D. W. (1988). *Principles of artificial intelligence and expert systems development.* New York: McGraw-Hill.

Waterman, D. A. (1986). *A guide to expert systems.* Reading, MA: Addison-Wesley.