Journal of Knowledge Management

Knowledge modeling for the preservation of institutional memory

John W. Coffey
Assistant Professor/Research Associate in the Department of Computer Science and the Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32051, USA
Robert R. Hoffman
Research Scientist at the Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32051, USA

Emerald

Journal of Knowledge Management, Vol. 7 No. 3, 2003,

Knowledge modeling for the preservation of institutional memory

John W. Coffey and Robert R. Hoffman

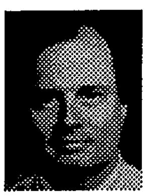


John W. Coffey is an Assistant Professor/Research Associate in the Department of Computer Science and the Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32051, USA (jcoffey@ufw.edu).
Robert R. Hoffman is a Research Scientist at the Institute for Human and Machine Cognition, University of West Florida, Pensacola, FL 32051, USA (rhoffman@ai.ufw.edu).

Abstract After setting the stage by briefly surveying knowledge elicitation techniques, this article presents a description of an iterative approach to the elicitation and representation of organizational knowledge called PreSERVe, which stands for prepare, scope, elicit, render, and verify. The method involves an initial process of preparing for knowledge elicitation, followed by an iterative process of assessing the scope of the endeavor, knowledge elicitation and rendering, and, verification. Use of the PreSERVe method is illustrated by a case study involving work with six senior engineers at NASA Glenn Research Center (NASA GRC), Cleveland, OH, USA.

Keywords Modelling, Knowledge management, Memory, Concept modeling

Introduction

Not long after the introduction of the seminal ideas of expert systems, a problem called the "knowledge acquisition bottleneck" was identified (Cullen and Bryman, 1988; Duda and Shortliffe, 1983) – it could take longer to elicit and represent the expert's knowledge than to actually program a rule-based system. What ensued were empirical evaluations of alternative knowledge elicitation (KE) methods (Burton *et al.*, 1987; Hoffman, 1987) and attempts to provide software support for the knowledge acquisition (KA) process (Boose, 1986). It was recognized that the acquisition of expert knowledge was a process having a number of steps, including knowledge elicitation, knowledge representation, implementation, and then validation or verification (Ford and Bradshaw, 1993; Regoczei and Hirst, 1992). Coinciding with these developments was a growing recognition that the technologies of KE and KA held promise for helping organizations cope with the loss of institutional knowledge (Klein, 1992; Seifert *et al.*, 1997). We now find ourselves in an era of rapidly shifting corporate and government workforces, and the preservation of institutional knowledge has become a widespread concern (Allee, 1997; Brooking, 1991; Choo, 1998; Davenport and Prusac, 1998; Von Krogh *et al.*, 2000). Institutional memory loss is a significant problem that can impact an organization's ability to advance its mission successfully, its ability to avoid making the same mistakes it made in the past, and its ability to leverage the accomplishments of departing employees.

"Institutional memory loss is a significant problem."

Various methods have been used to attempt to preserve institutional memory. Knowledge preservation (KP) techniques can borrow from KE and KA activities of a variety of types. Each of the methods has strengths and weaknesses. Weaknesses often take the form of difficulties that stem from the tradeoffs between the ease of acquisition and the usability of that which is acquired. Here are four examples:

(1) Knowledge acquisition methods can consume large amounts of time of busy experts, and take them away from their main tasks. It is difficult to convince time-pressed employees to record information in an ongoing fashion, which mediates against the effectiveness of approaches that require ongoing collection of information.

(2) Exit interviews may yield useful information, but they are more likely to be brief and superficial or unworkable if the employee is exiting under less than desirable circumstances. When more exhaustive KE is attempted, as in the method of "oral histories" (Paris, 2001) large quantities of information can be acquired, but such information may not be in an easily used form (i.e. it may be extended text or video), making the needed information difficult and time-consuming to access.

(3) The search for information in large corporate archives is difficult for a host of reasons, many having to do with deficiencies in indexing. Attempts to automate this process have helped to spawn an entire field called data mining (Chua *et al.*, 2000; Holmes *et al.*, 2000; Smith *et al.*, 2000).

This paper presents the PreSERVe (prepare, scope, elicit, render, verify) method of knowledge modeling. It is an iterative approach to knowledge modeling in service of KP. The method involves an initial process of preparing for knowledge elicitation, followed by an iterative process of assessing the scope of the endeavor, eliciting and rendering knowledge by various methods, and verifying the renderings with the experts. This approach was developed and refined in work at NASA Glenn Research Center. The efficacy of any such method may be judged by a variety of parameters including acquirability, accessibility, usefulness, and usability. As we will show, the PreSERVe method affords a reasonably high level of acquirability, easy accessibility to the acquired knowledge, yields a very high proportion of potentially useful knowledge, and preserves the knowledge in a form that is useable.

This article begins with a brief survey of KE and knowledge representation techniques, followed by a description of the PreSERVe method, and then the case study in its application to the field of "launch vehicle systems integration" at NASA GRC.

Knowledge elicitation techniques

Approaches to KE have been defined primarily by work in the area of expert systems (see Waterman, 1986). However, attempts to preserve institutional memory have a broader focus than the task of creating an expert system. Approaches to KE have been defined by work in the psychology of expertise and cognitive anthropology as well as the work in expert systems (see Hoffman, 1992; Hutchins, 1995). Among the many dozens of KE methods are the critical decision method (Klein *et al.*, 1989; Crandell and Getchell-Reiter, 1993), the knowledge audit (Hoffman *et al.*, 2000), the cognitive modeling procedure (Hoffman *et al.*, 2000), protocol analysis (Ericsson and Simon, 1993), work patterns analysis (Vicente, 1999), oral historiography (Paris, 2001), exit interviewing (Jurkiewicz *et al.*, 2001; Kransdorrf, 1997), and active and retroactive archiving (Crowley, 1997; Griffith, 1997). Hoffman *et al.* (1995) categorized KE methods as involving:

- analysis of the tasks experts usually perform as they engage in problem solving or decision-making (task analysis, work flow analysis, workspace analysis, etc.);
- interview techniques (unstructured, or structured in any of a variety of ways); and
- contrived techniques that bring to bear some sort of experimental or semi-experimental manipulation of the expert's familiar tasks.

The following brief discussion describes these general categories in a bit more detail.

The analysis of familiar tasks

This family of methods encompasses a range of activities that involve observation of the expert actually performing work. In this approach, the expert is sometimes asked to verbalize, permitting the analysis of a protocol for the task. This method also can utilize test cases that may be typical of the sort of problems the expert faces, or that may be anomalous or difficult.

Interviews

Unstructured interviews can be used to achieve rapport, and to gain a global overview of the domain under consideration. However, some degree of structure is generally preferable for all but the initial interviews. Structured interview methods come in dozens of forms, but all of them are organized around what Hoffman *et al.* characterize as "domain-specific probe questions".

Contrived techniques

This category includes such methods as decision analysis, rating and sorting tasks, graph construction, and constrained processing/limited information problems. In decision analysis, the expert generates lists that include the elements of a problem, relationships among the elements, the types of problems encountered, etc. From such an analysis, the approach to decision-making can be culled (i.e. a reasoning model). Rating and sorting tasks involve making judgments regarding attributes of a problem domain (e.g. the relations among domain concepts). Graph construction involves the creation of conceptual graphs or concept maps (Novak and Gowin, 1984) that are diagrammatic representations of knowledge. Concept maps are diagrams that are comprised of concepts in the nodes, and linking phrases that elaborate the relationships among concepts on the arcs. The elicitation of concept maps has proven to be an effective means of externalizing an expert's key concepts of a knowledge,domain (Ford *et al.*, 1993; Ford and Bradshaw, 1995). Figure 1 presents a concept map that is about concept maps. Concept maps are useful for the determination of the scope of a discussion by elaborating the defining, basic concepts to be considered. Interview structure can be based on

Figure 1 A concept map about concept maps (adapted from Novak and Gowin, 1984)

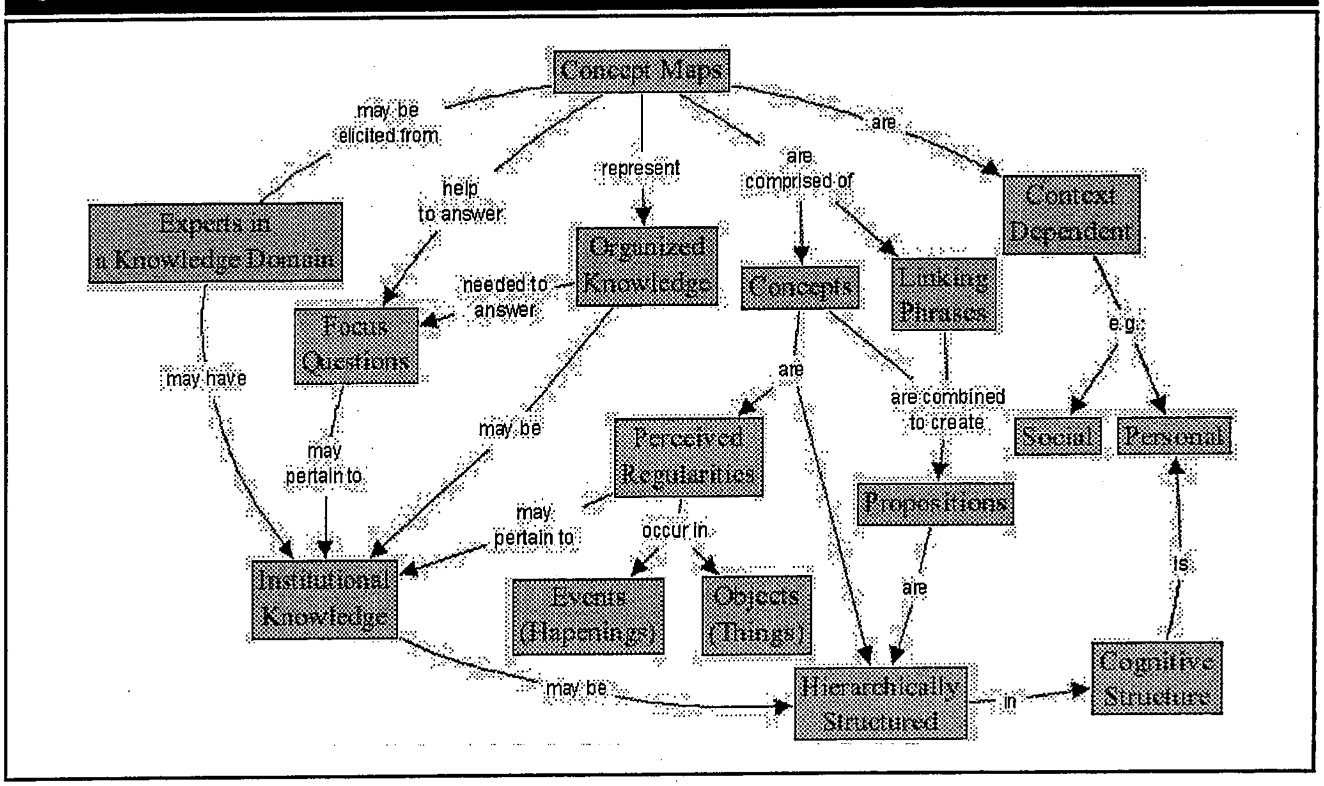

the elicited concepts in the concept maps. The act of working through the issues in the context of the concepts from the concept maps creates an efficient framework for discussions.

The various KE techniques can be used in different combinations and with a focus on eliciting knowledge of a certain type in order to serve a specific purpose. For example, the critical decision method (Klein *et al.*, 1989; Crandell and Getchell-Reiter, 1993) attempts to capture "naturalistic decision making". The expert is instructed to recount a specific problem or case that was especially difficult, or in which the expert made a decision that went against the generally held view in a way that ultimately proved to be correct. The approach combines structured interviews, protocol analysis, case-based reasoning, and retrospection. Another hybrid method of knowledge acquisition is based upon concept maps and repertory grids (Cañas *et al.*, 1994; Coffey, 1995; Ford *et al.*, 1991, 1993, 1996).

It is the notion of combining KE methods that forms a key aspect of the PreSERVe method, because, as we will show, the method takes the results from KE and puts them in a broader framework in which knowledge is preserved in a form and format that is potentially useful and useable. The basic goals of KE for KP are taken to be: eliciting knowledge efficiently, efficiently converting the knowledge into a concise electronic form, organizing the knowledge in a way that makes it accessible, and verifying the correctness of that which has been acquired. The PreSERVe method creates a framework in which to achieve these goals.

The PreSERVe method

Figure 2 presents an overview of the PreSERVe method. In the sub-sections that follow we discuss each of the components of the method in order.

Preparation

Preparation has several constituents including the identification of the expert or experts (if necessary), the selection of the knowledge domain to be explored, and advance preparation by the knowledge engineer(s) who will facilitate the process. Preparation may be carried out in a variety of ways. Alternatives include independent reading of materials identified by the

Figure 2 The PreSERVe method

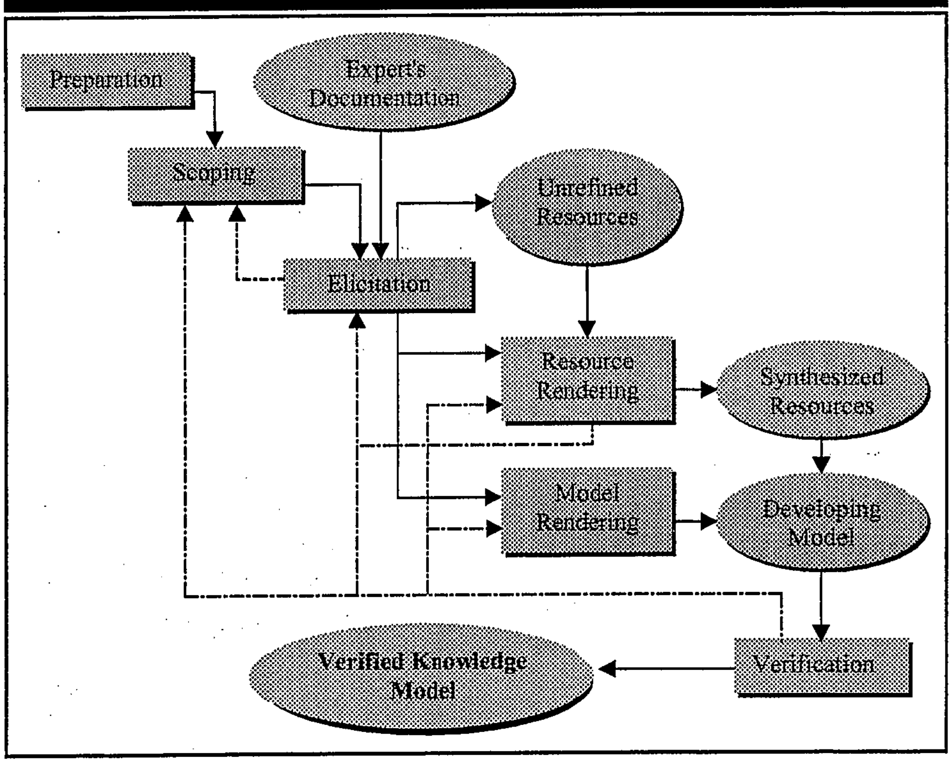

knowledge engineer or provided by the expert, perusal of company records and other documentation, preliminary, informal meetings with principal participants, etc. This component has been described as "bootstrapping" in the literature on KE (Hoffman *et al.*, 2000, 1998). Another aspect of the preparation step is proficiency scaling – making the best determination possible of the true level of expertise in those who might participate. Proficiency scaling makes it possible to establish criteria for selecting the participants from whom knowledge will be elicited. Once the participants are selected, another aspect of preparation that is critical is the need for the knowledge engineer to establish rapport with the expert(s). The expert's attitude toward the effort plays a significant role in the success or failure of the endeavor.

Scoping

It is important to make a clear determination of the scope of the KE and KP endeavor. In the realm of KP, it is to be expected that the scope might broaden or narrow as the project unfolds. The need to maintain and act upon an awareness of the potentially changing scope and direction of a KP project is one of the major driving forces in defining the PreSERVe method as an iterative process. The development of focus questions helps maintain direction (Novak and Gowin, 1984). An overriding concern in the creation of focus questions is to formulate ones that are at the right level of granularity – neither too general nor too specific. A failure to establish the scope of a knowledge elicitation project can lead to a haphazard elicitation that may go in directions that do not serve the goal of the project, and that may neglect areas that are important to the goals. For this reason, it is ideal (but very difficult) to set the scope of the project so that all the significant issues can be addressed with the allotted resources. A more feasible policy is to identify topics as carefully as possible, to rank order the topics by importance (in conjunction with the experts) and to elicit all the topics that time allows.

Elicitation

Having determined a scope that is reflected in the focus questions, the actual elicitation of knowledge commences. Methods are chosen from those enumerated in section 2, typically with concept mapping playing a central role. Often, in the course of the KE, experts will refer to documents or other resources (characterized as unrefined resources in the PreSERVe method) to illustrate or elaborate upon their ideas. These resources include those that are culled from the expert's accumulated documentation, and can take a variety of forms such as hardcopy or electronic documentation upon which the expert relies, audio or video of the expert that may be edited, etc. This documentation is both an input to the knowledge elicitation process and a source of materials that can be used as auxiliary knowledge resources that elaborate the knowledge model.

Rendering

Rendering takes two forms: the creation of elements that will be included in a knowledge model (resource rendering), and the assembly of these elements into the knowledge model itself (model rendering). The KE process leads to the creation of artifacts (concept maps, interview transcriptions, edited videos of the expert discussing a topic or making a point, etc.) that are included in the knowledge model. These and other pre-existing resources that have been identified in the process (such as selected items from the expert's reference material) are rendered in a form that is suitable for inclusion in the knowledge model. The totality of these edited resources are characterized as synthesized resources. These resources must have an organizing scheme that enables the model builder to organize all the resources in a logical way, and to identify redundancies and gaps. The knowledge model may be structured by the concept maps themselves or by any other organizing scheme that can be utilized to build hypermedia.

Verification

Verification of a knowledge model must occur at both a conceptual and mechanical level. A typical online knowledge model is comprised of many sources of information linked together into a cohesive whole. A conceptual or semantic verification (the more critical of the two forms) is the process of checking that the factual information is correct, that the best content or information available at the time is included, that the overall structure, arrangement and retrieval mechanisms make sense, and that errors and significant redundancies have been eliminated. Ascertaining that all the components can be found, that the links work, and that the media display correctly, corresponds to a mechanical verification. Verification may be performed by independent experts. However, the less certainty there is in a knowledge domain, the less useful independent attempts at corroboration will be.

The iterative nature of the method

The PreSERVe method is iterative. Broken lines in Figure 2 illustrate potential iterative loops. After a round of KE, it is possible that the scope of the undertaking will change as the experts reflect on the topics that have been elicited and their relative importance. An evaluation of the scope of the project can be undertaken after each round. It can become apparent that the scope of the undertaking is too broad or that critical elements have been omitted. An assessment of the scope must be made to ensure that the limits of the project are both comprehensive and can be met, and that the actual work continues to fall within what is planned. After the rendering work, a review of the model with the expert may lead to the conclusion that additional resources should be included (a revisit to the resource rendering phase), that more knowledge should be elicited (revisiting the elicitation phase), or that the scope should be broadened or narrowed.

The number of iterations is determined by several factors, including the size of the undertaking, the skill of the knowledge engineer and support personnel, and the time allotted. Eventually, a deliverable version of the knowledge model is produced. If the knowledge elicited is essentially historic, the model will remain stable except for possible additions due to subsequent identification of historical content. If the domain is dynamic and changeable, knowledge model maintenance becomes a critical concern, an issue that is outside the scope of this work. The next section describes a case study in the use of this method.

A case study – knowledge preservation at Glenn Research Center

The PreSERVe approach was used in order to capture, represent, and preserve institutional memory of senior scientists at NASA's Glenn Research Center, regarding issues pertaining to launch vehicle systems integration (Coffey *et al.*, 1999). The Centaur upper stage launch vehicle and the RL-10 engine were chosen for detailed analysis. A total of six engineers, all of whom were at or near retirement from NASA or NASA contractors, were involved in this project. This case study will first be described as it unfolded chronologically. Following this discussion, details of knowledge elicitation and representation will be described.

Chronological account

The preparation phase proceeded with the knowledge engineer reading several tracts on thermo dynamics, fluid mechanics, the basic design of liquid propellant rocket engines, and appropriate chapters of a tract on launch vehicle design by Husal and Huang (1992). These sources were obtained on the basis of recommendations from the expert and from initiatives by the knowledge engineer.

The initial goal of the project was to create a high-level description of the concerns of an engine designer as they contrast with those of a vehicle designer. This description of the scope of the project was subsequently deemed to be excessively broad. At an early meeting, the decision was made to focus on vehicle and engine integration issues as they pertain to the Centaur/ RL-10 upper stage. The basic focus question that defined the initial sessions was: "What are the major concerns of the engine and vehicle designers of the Centaur upper stage and RL-10 engine, and in what ways are these concerns complementary or contradictory?". Figure 3 presents a concept map that was created to guide the initial KE effort.

The first round of KE

Three trips were made to the NASA Glenn Research Center. Each trip had a duration of three working days. Initially, the expert laid critical groundwork for the knowledge acquisition team by

Figure 3 A concept map used to guide the initial KE sessions

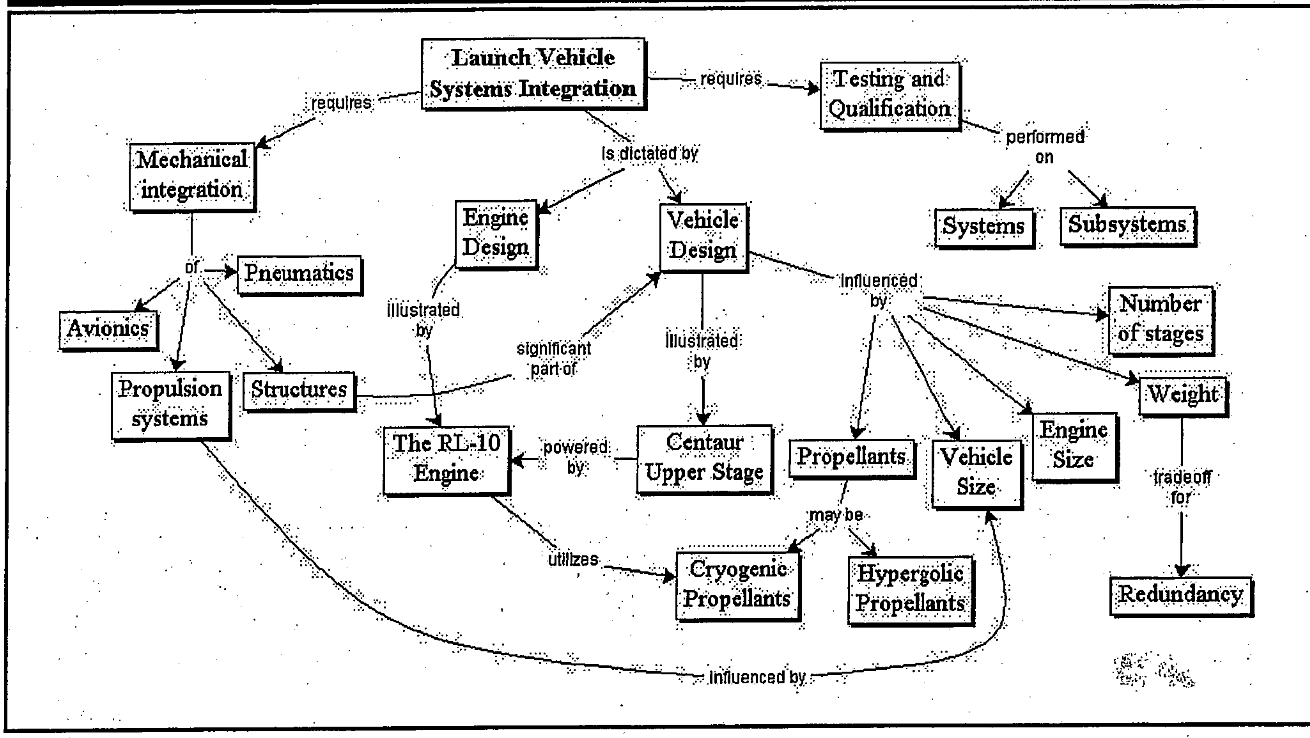

helping the team to understand his unique perspectives on the many engineering issues pertaining to cryogenics, propellants, expander cycle engines, the evolution of the RL-10 and the Centaur vehicle.

In the times between the visits, work on the rendering of the knowledge model proceeded. The rendered elements were created and then reviewed with the expert on each subsequent visit. Rendering work on the first version of the knowledge model involved the creation of text passages from the notes that were taken in the interviews, and initial attempts to match them to concepts in the concept maps. Excerpts of digital videos of the expert discussing points of interest were created. The notes and concept maps were then verified by the expert. The final result of this part of the effort was a first version of a browsable, hypermedia model of the expert's domain knowledge of the Centaur/RL-10 system.

The second round of KE

At the end of the first round, the expert suggested that a review of the various failures that occurred in the Centaur/RL-10 program would uncover significant issues pertaining to the integration of launch vehicle and engine. This insight led to a second adjustment in the scope of the project and the format of knowledge elicitation, in order to accommodate the elicitation of knowledge from multiple experts. The result was a plan to convene a group discussion of the various failures in the Centaur program. This group discussion was called the "Failure Panel". The Panel was comprised of a group of five experts in the Centaur/RL-10 system who discussed the various mission failures in the Centaur program and the lessons learned in those failures.

The discussion spanned two days, and developed into an extensive discourse on design, testing and integration of subsystems and systems, and NASA's oversight role in these activities. Issues such as the design, fabrication and assembly of components, subsystems and systems, and issues such as getting the engines and vehicles to Cape Canaveral, production and storage issues, etc. were addressed. The sessions were audio and video taped. In the course of the discussions, many relevant engineering drawings were identified and retrieved from the NASA archives. These drawings were part of a very large collection. It was really quite difficult to determine the significance of any individual drawing until it became relevant to these discussions. The experts helped to identify engineering drawings that illustrated the critical components that were of interest and that merited discussion and elaboration. The elicited knowledge from this panel was transcribed and rendered into Web pages with illustrative complementary media (graphics and videos) interspersed in the texts.

The result of the Failure Panel was a multifaceted perspective on launch systems integration, and particularly how it could fail, from the people who worked on integrating the onboard computer, navigation and guidance systems, propellant tanks, piping and the engines. As the Panel proceeded, open issues that could not be resolved by the panel were recorded. A review of the issues revealed that many of them pertained to particulars regarding the design of the RL-10 engine itself.

The third round of KE

The Panel identified a recent Pratt and Whitney (manufacturer of the RL-10) retiree, who currently worked for a sub-contractor. This expert was consulted regarding open issues that the Panel had enumerated relative to the RL-10. The first day of work with this expert was comprised of discussions of the open issues from the Failure Panel, presentations by the expert, and questions from the knowledge engineers. Day 2 consisted of structured interviews based upon unresolved issues from Day 1. The knowledge that emerged from these sessions was interesting because a new perspective was added to the mix – that of a prime contractor who had worked with NASA and third-tier subcontractors to design, build and test engines. Hence, this expert was in a position to describe the process of resolving the conflicting priorities of the engine and vehicle maker from the engine manufacturer's point of view.

A total of five multi-day KE waves were involved in this project. The first three were with the original expert, the fourth with the Failure Panel, and the fifth with the engineer from Pratt and Whitney. Verification of the elicited materials commenced as soon as the materials were rendered after each round. Content changes were typically minor, and the construction of the knowledge model was done in consultation with the experts involved. At the conclusion of the effort, the model was delivered to NASA to be utilized for training and education. The following sub-sections elaborate the methods of knowledge elicitation and representation that were used in these rounds of KE.

Knowledge elicitation methods

A variety of KE methods were used in this process, with different approaches selected to match the phase of knowledge elicitation, whether individual or multiple experts were involved, and the basic goals of the session. KE in the first round typically started with a contrived task, the creation of concept maps (Novak and Gowin, 1984) for the domain in question. Starting with the concepts in the concept maps led to extended discussions that were augmented and made comprehensible by incorporation of secondary sources of information that the expert identified. The Failure Panel was structured by several principles:

- the failures would be considered chronologically;
- each Panel discussion was moderated by the expert who was most familiar with the area in which the failure occurred; and
- the focus of the discussion was the cause of the failure and lessons learned for the engine/vehicle integration process.

It was agreed in advance that if a failure provided no insights regarding engine/vehicle issues, or if it was in an area that was outside the scope of the experts' expertise, that it would not be considered. Only a few were discarded. The Panel discussion ranged over matters of engine and vehicle design, systems integration, the exigencies of reliable fabrication, and many other topics. All of these sessions were audio and video taped and transcribed. A propositional analysis of the discussion notes was performed to determine the major issues that were revealed by the discussions. These became entries in one or more of the final model's indexes.

The final KE sessions were with the engine designer. This expert had a clear idea of what he wanted to make known regarding his experiences in engine development and testing. One of the most interesting outcomes from these sessions was the creation of a set of principles for the design and testing of expander cycle engines. These were very broad design principles, backed with significant detail from the interviews. These principles had always been implicit in the expert's thinking, but had never been formally organized as they were as a result of the KE process.

In summary, the knowledge elicitation methods encompassed all three categories of methods described above. Unstructured interviews were valuable in the initial task of attempting to decide how to characterize the expertise of the expert from whom knowledge would be elicited. The contrived method of creating concept maps to guide interviews was successfully employed in attempts to determine the scope of discussions and as a yardstick to judge after the fact, where gaps in the elicited knowledge existed. Structured interviews that were taped and in some cases transcribed (a labor-intensive process), added additional substance to the skeletal outlines created with the concept maps. Finally, the analysis of familiar tasks was very useful as experts considered reasoning strategies that led to system designs, the determination of failure causes, and the resolution of design issues with conflicting requirements.

This KE effort encompassed a total of 17 working days of contact with the experts. This was a brief time commitment in terms of face-to-face contact with the experts themselves. However, significantly more time was required for the knowledge engineers and support personnel in preparation, review and transcription of notes, for the creation of the knowledge model, etc. Additionally, substantial effort was expended by the experts off-line to review the emerging knowledge models. However, since the experts had worked for NASA or contractors for an average of more than 30 years, the result was a corpus of knowledge that included many of the highlights and lessons learned in 180 person/years of engineering work.

Model creation, organization and user interface

As the rounds of KE progressed, three separate models were developed. The original knowledge model, structured around the concept maps that were elicited from the original expert, was crafted from interviews and the resources he identified. The Failure Panel discussions were divided into chapters for each of the individual failures that were discussed. The transcripts of the sessions were rendered in web pages, with the engineering drawings to which the experts referred interspersed in the text. The work was indexed in several ways: chronologically, by failure, by individual topics within a given failure, etc. A third model was created from interviews with the engine designer. This model was organized around the basic principles for design of expander cycle engines that had been identified in the knowledge elicitation effort.

The rendering process initially proceeded separately on the three separate knowledge models. When it had been determined, in review with the experts, that a fairly comprehensive rendering of each of the three components had been completed, all three knowledge models were integrated together by cross referencing key concepts so that the various media were accessible through the various indexes.

A total of 15 concept maps, 118 graphics, 102 textual passages, and 18 digital videos were included in the combined knowledge model. The graphics were all identified in the course of the knowledge elicitation sessions. The areas for which videos were created were identified during the knowledge elicitation sessions as those that would best benefit from such a presentation.

A global indexing scheme was created for the basic concepts in the integrated model. Figure 4 contains a screen shot of the basic indexing schemes for the system. The backmost window labeled "Launch vehicle systems integration" is the front page to the entire model. From that page, the user may navigate directly into the model by a variety of means, search for information of interest, or select topics of interest from the indexes.

The user can navigate directly to any of the three individual models by clicking on the appropriate name under the "As told by" prompt on the front page. The user can select to view a global index of the entire integrated model by clicking the "Global index" label that also

Figure 4 Indexing schemes in the NASA Glenn Research Center knowledge model

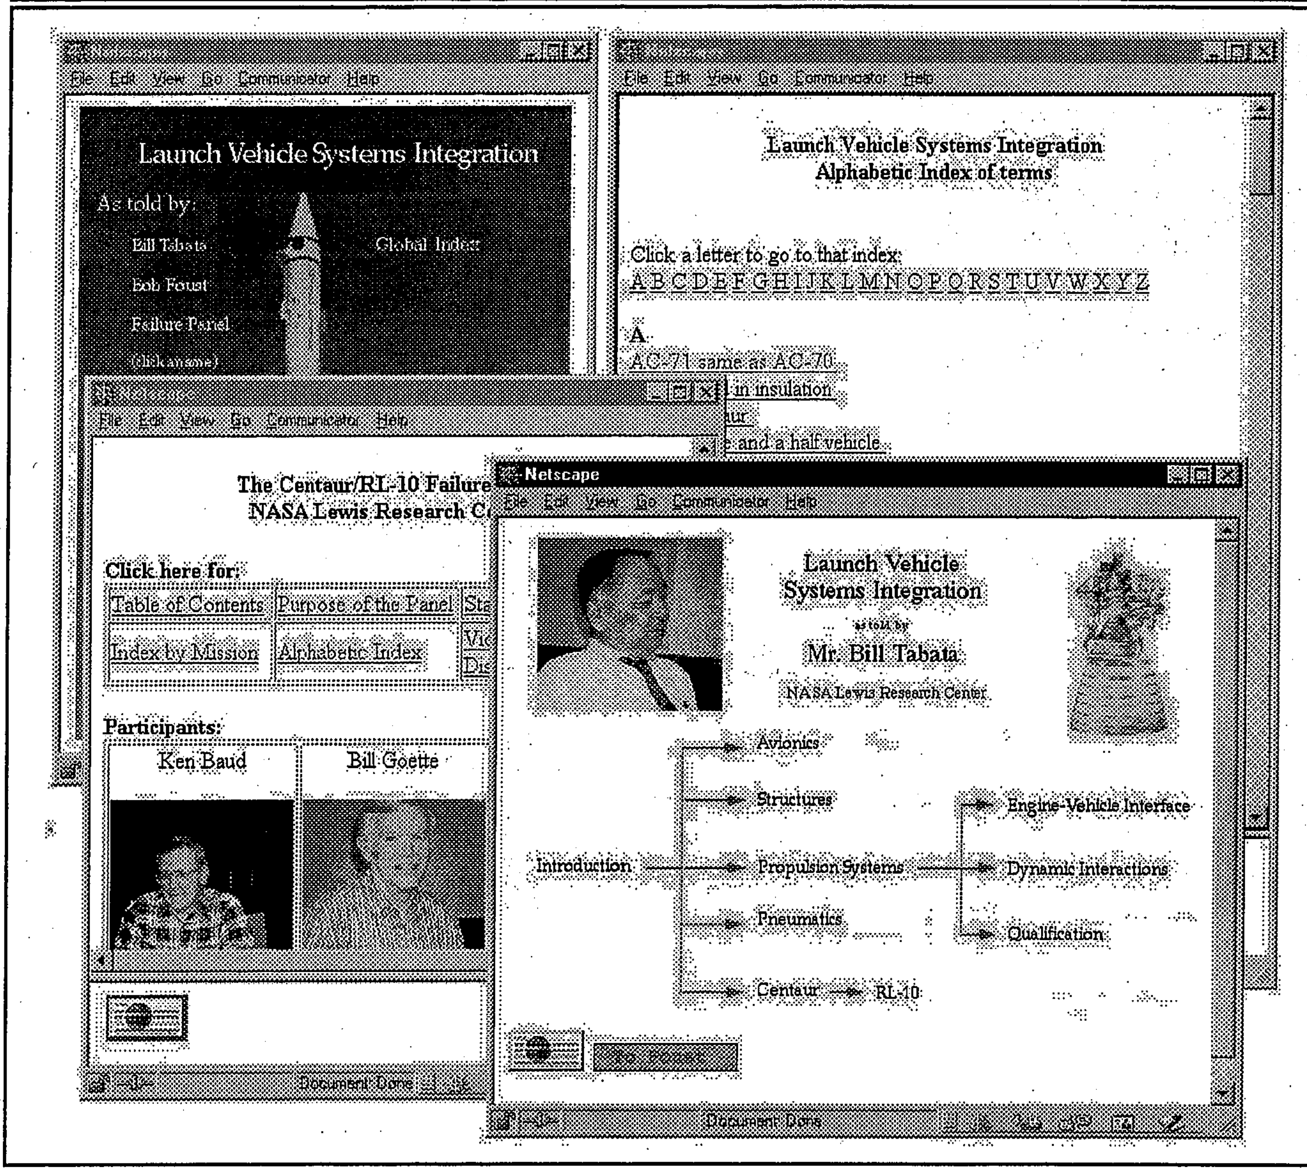

appears on the front page. That action would present the "Alphabetic index of terms" window that is visible as the top right window in Figure 4. The individual models have similar indexes that structure the topics as shown in the bottom-right window in Figure 4. This window illustrates how these indexes present a hierarchical ordering, starting with a global, introductory view of the model. The most general view can be accessed by clicking on the label that reads "Introduction". Beneath the introduction is the next level of detail, avionics, structures, propulsion systems, pneumatics, and Centaur/RL-10. Further detail is presented on propulsion systems including engine-vehicle interface, dynamic interactions and qualification issues regarding components, subsystems and systems.

Figure 5 depicts a typical navigation path that a user might follow through the model. In this scenario, the user viewed the propellant flow schematic of the RL-10A-3-3A. From that window, the user could have gone to the index, to a text on the subject of the RL-10, to other graphics of the RL-10, to a video, or to several concept maps that contain information on the RL-10. These possible locations are indicated by pull-down menus associated with the icons at

Figure 5 A typical navigation sequence in a knowledge model of the sort created with the PreSERVe method

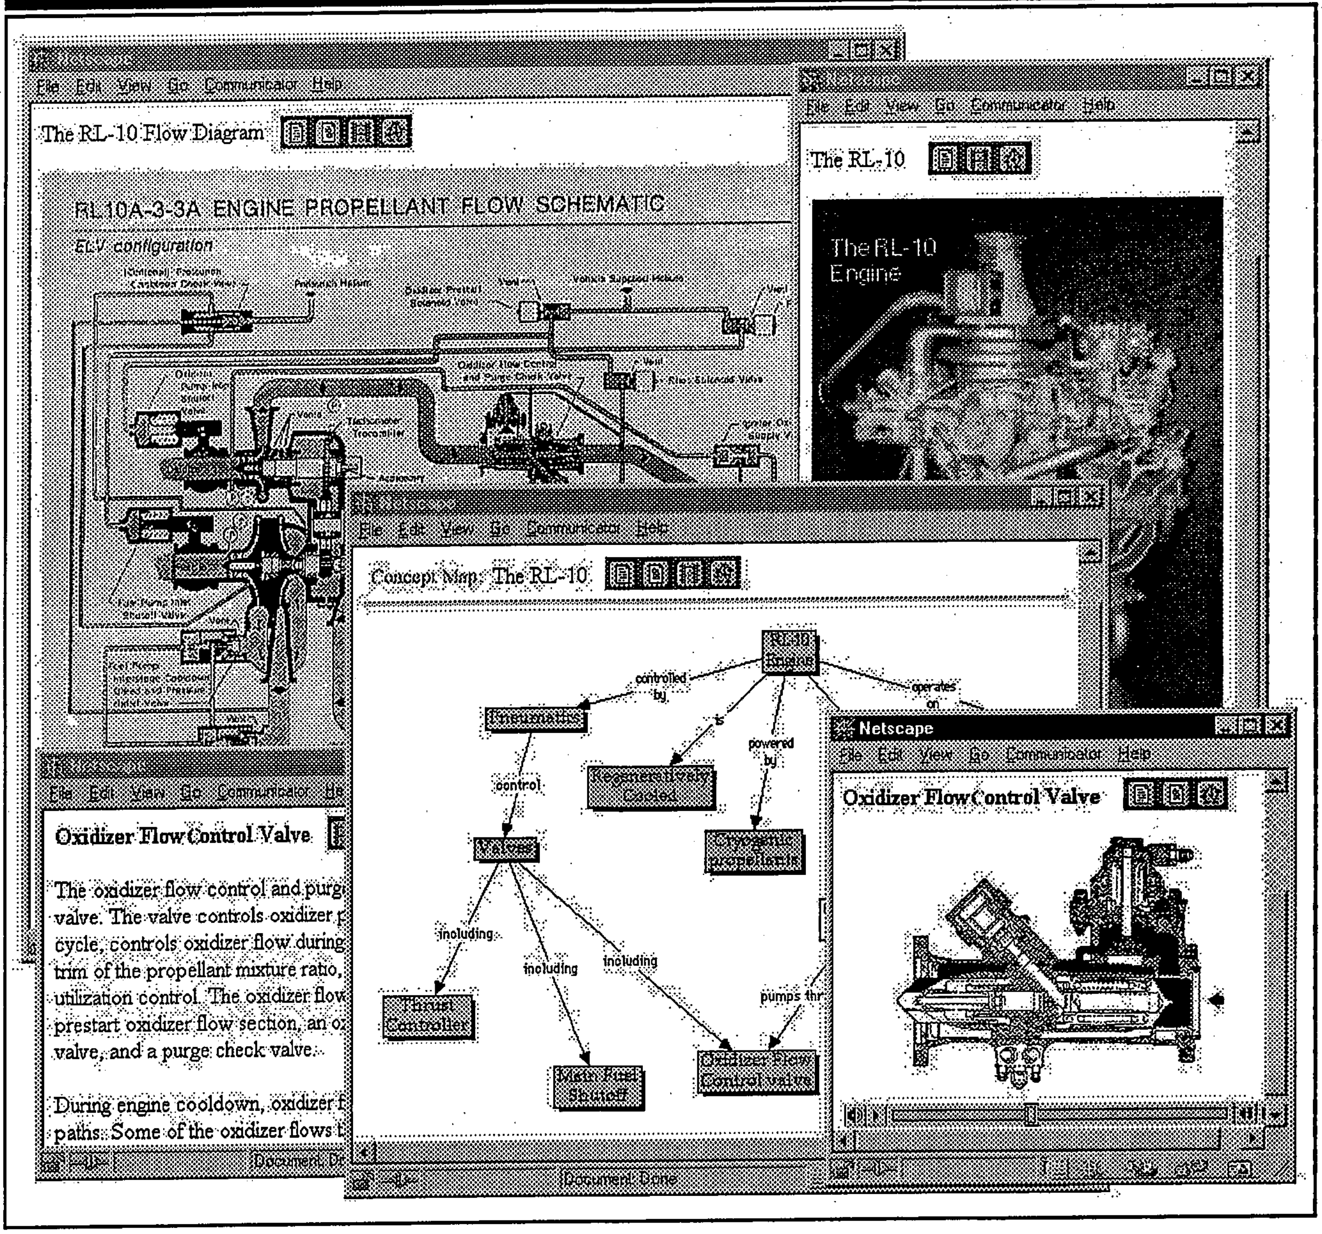

the top of each window. From left to right in the RL-10 flow diagram window, the icons represent textual passages, graphics, videos and Concept Maps. From the propellant flow schematic, the user navigated to a graphic of the RL-10, to a concept map that was used to organize the discussions on the RL-10, to a textual passage concerning the RL-10's oxidizer flow control valve, and then to a video of the expert describing design issues in the context of the oxidizer flow control valve. The knowledge model presents a multimedia knowledge space that can be freely browsed in the manner described above, or that can be directly accessed from any of the various indexes illustrated in Figure 4.

Summary

The problem of preventing the loss of expert knowledge – knowledge preservation – is pervasive today and will only worsen as the post-war "baby boom" generation approaches retirement age. Knowledge elicitation and modeling can be combined to serve as one in a range of approaches to address this problem. This paper presents a description of the PreSERVe method (prepare, scope, elicit, render, andverify) of knowledge modeling. This method starts with an initial phase of preparing for a knowledge modeling effort, followed by an iterative process of examining scope, eliciting knowledge, rendering that knowledge in a computerized form, and verifying the knowledge with the expert. This work surveys KA techniques that are relevant to the PreSERVe method, describes the PreSERVe method itself, and presents a case study in its use at NASA Glenn Research Center.

Discussion

The role of concept maps and other KE methods

The use of concept maps to gain a global perspective helps ameliorate the problem of knowing when a comprehensive account has been created. An approach that starts with the elicitation of concept maps allows the expert to define the boundaries of the knowledge domain – to assess the current scope of the project. After a short period of map building followed by a review of the map, the expert can more readily recognize if all the key ideas have been included. Interestingly, the approach of working on a concept map, taking some time away from it, and then revisiting the map is a highly effective, iterative way to ensure comprehensive coverage. Map creators invariably want to change things upon re-examination. The structure of a well-constructed concept map presents a clear indication of the relative importance of the concepts. The concept maps can be retained and used as the organizing structure for the knowledge model itself.

The PreSERVe method provides a framework for the utilization of the full range of KE approaches. The concept maps provided very concise snapshots of the items to be considered. Unstructured interviews were useful initially in determining the goals of the project and in deciding on the need for iterations of various activities. Extensive structured interviews were conducted. The use of concept maps was a critical factor in avoiding a biasing of the sessions in the more structured interviews. The analysis of familiar tasks played an extensive role in the analyses of how the failures were examined and lessons learned determined.

Relative efficiencies in the PreSERVe method

It is critical to the success of KP methods such as PreSERVe that knowledge can be elicited efficiently, converted into a concise electronic form, organized in a way that makes it accessible, and verified for correctness. The PreSERVe method creates a framework in which to achieve these goals. Efficiencies in the knowledge elicitation effort are attained because the method keeps the work well focused on the issues at hand. The parallel work on resource identification and basic resource rendering creates efficiencies in model construction. As resources are identified, they can be incorporated into the knowledge model. The ongoing review of the emerging model efficiently identifies areas that need further consideration with the experts in additional rounds of knowledge elicitation or resource identification. The emphasis on the ongoing re-evaluation of scope is critical since it is highly improbable that all the important elements of a group of experts' experiences can be anticipated in advance.

Although the PreSERVe method presents a well-specified method that seeks to exploit efficiencies where it can in pursuit of KP, it is acknowledged that this approach places the quality and utility of the final result above the efficiency of the knowledge elicitation process. The elicitation of expert knowledge seems highly resistant to significant automation. For this reason, substantial human effort is necessary for its preservation. However, a method such as the PreSERVe method affords a principled approach. Further refinement of the method and additional work on knowledge model maintenance will take place in future work.

“ The elicitation of expert knowledge seems highly resistant to significant automation. ”

> “They were still excited about the opportunity to discuss what they knew and what they had learned.”

Lessons learned in the case study

The iterative nature of the PreSERVe work served the process well. The project encompassed three rounds of knowledge elicitation with a total of five multi-day waves. The need for and focus of subsequent waves was identified as the work unfolded. This fact clearly supports the need for the careful assessment of scope as an ongoing part of the process.

The incremental approach of verifying artifacts as they were produced and organized served the effort well. The separate development of the three components of the knowledge model and their subsequent integration is represented explicitly in the model's render-verify loop. The iterative approach also proved useful in helping to uncover additional resources that contributed to the knowledge model. The PreSERVe method's support for the identification of accompanying media, (graphics, video, etc) is another strength of the approach. The engineering drawings, literal images, selected videos of the experts, etc., contribute substantially to the richness of the final knowledge model.

The ongoing review and evaluation explicit in the PreSERVe method helped the experts to identify gaps in their own knowledge and then identify other experts who could fill those gaps. Much of the knowledge that was elicited would be described as tacit – the idiosyncratic understandings of the experts. A reference such as Huzel and Huang is a valuable resource, but based upon well-known principles. This effort resulted in a much greater proportion of the more unique, tacit knowledge of the experts – the kind of knowledge that is the most difficult to preserve.

One unanticipated result of this effort was the discovery of how willing the experts were to discuss their work. These folks had worked in the area of launch vehicles for their entire careers and they were still excited about the opportunity to discuss what they knew and what they had learned. Their knowledge was available for the capture; it was only an issue of ensuring that someone would take the initiative to elicit it in a principled way. The PreSERVe method presents such a principled approach, and as a consequence, many of the major lessons learned in 180 person-years of work on launch vehicles and engines were retained.

References

Allee, V. (1997), *The Knowledge Evolution: Expanding Organizational Intelligence*, Butterworth-Heinemann, Boston, MA.

Boose, J.H. (1986), *Expertise Transfer for Expert System Design*, Elsevier, Amsterdam.

Brooking, A. (1991), *Corporate Memory: Strategies for Knowledge Management*, Thomson International, London.

Burton, A.M., Shadbolt, N.R., Hedgecock, A.P. and Rugg, G. (1987), "A formal evaluation of knowledge elicitation techniques for expert systems: domain 1", in Moralee, D.S. (Ed.), *Research and Development in Expert Systems*, vol 4, University Press, Cambridge, pp. 35-46.

Cañas, A.J., Ford, K.M. and Coffey, J.W. (1994), "Concept maps as a hypermedia navigational tool", Presentation at FLAIRS-94, The Florida AI Research Symposium, Pensacola Beach, FL, 6 May.

Choo, C.W. (1998), *The Knowing Organization: How Organizations use Information to Construct Meaning, Create Knowledge, and Make Decisions*, Oxford University Press, New York, NY.

Chua, C., Chiang, R., Lim, H.-L. and Ee-Peng, L. (2000), "Integrated data mining system to automate discovery of measures of association", Presentation at The 33rd Annual Hawaii International Conference on System Sciences (HICSS-33), Maui, USA, 4-7 January.

Coffey, J.W. (1995), "Issues in hypermedia and participatory explanation", *Proceedings of the Eighth Florida AI Research Symposium (FLAIRS '95)*, Melbourne, FL.

Coffey, J.W., Moreman, D. and Dyer, J. (1999), "Knowledge preservation at Glenn Research Center", Technical Report, Glenn Research Center, Cleveland, OH.

Crandell, B. and Getchell-Reiter, K. (1993), "Critical decision method: a technique for eliciting concrete assessment indicators from the 'intuition' of NICU nurses", *Advances in Nursing Sciences*, Vol. 16, pp. 42-51.

Crowley, X. (1997), "Memory bank: companies try to record the knowledge of their information workers", *PC Week*, Vol. 14, pp. 101-2.

Cullen, J. and Bryman, A. (1988), "The knowledge acquisition bottleneck: time for reassessment?", *Expert Systems*, Vol. 5, pp. 216-25.

Davenport, T.H. and Prusac, L. (1998), *Working Knowledge: How Organizations Manage what they Know*, Harvard Business School Press, Boston, MA.

Duda, R.O. and Shortliffe, E.H. (1983), "Expert systems research", *Science*, Vol. 220, pp. 261-8.

Ericsson, K.A. and Simon, H.A. (1993), *Protocol Analysis: Verbal Reports as Data*, MIT Press, Cambridge, MA.

Ford, K.M. and Bradshaw, J.M. (1995), "Beyond the repertory grid: new approaches to constructivist knowledge acquisition tool development", *International Journal of Intelligent Systems*, Vol. 8, pp. 287-333.

Ford, K.M. and Bradshaw, J.M. (1995), *Knowledge Acquisition as Modeling*, John Wiley, New York, NY.

Ford, K.M., Cañas, A.J. and Coffey, J.W. (1993), "Participatory explanation", *Proceedings of 6th Annual Florida AI Research Symposium*, Ft. Lauderdale, FL.

Ford, K.M., Cañas, A.J., Jones, J.C., Stahl, H., Novak, J.D. and Adams-Webber, J.A. (1991), "ICONKAT: an integrated constructivist knowledge acquisition tool", *Knowledge Acquisition Journal*, Vol. 3, pp. 215-36.

Ford, K.M., Coffey, J.W., Cañas, A.J., Turner, C.W. and Andrews, E.J. (1996), "Diagnosis and explanation by a nuclear cardiology expert system", *International Journal of Expert Systems*, Vol. 9, pp. 499-506.

Griffith, V. (1997), "Treasures in the corporate memory: usage of corporate archives", *The Financial Times*, 10 November, p. 14.

Hoffman, R.R. (1987), "The problem of extracting the knowledge of experts from the perspective of experimental psychology", *The AI Magazine*, Vol. 8, Summer, pp. 53-66.

Hoffman, R.R. (Ed.) (1992), *The Psychology of Expertise: Cognitive Research and Empirical AI*, Erlbaum, Mahwah, NJ.

Hoffman, R.R., Crandall, B. and Shadbolt, N. (1998), "A case study in cognitive task analysis methodology: the critical decision method for the elicitation of expert knowledge", *Human Factors*, Vol. 40, pp. 254-76.

Hoffman, R.R., Coffey, J.W. and Ford, K.M. (2000), "A case study in the research paradigm of human-centered computing: local expertise in weather forecasting", Technical Report, National Technology Alliance, Washington, DC.

Hoffman, R.R., Coffey, J.W. and Carnot, M.J. (2000), "Is there a 'fast track' into the black box?: the cognitive models procedure", Poster presented at the 41st annual meeting of the Psychometrics society, New Orleans, LA.

Hoffman, R.R., Shadbolt, N.R., Burton, A.M. and Klein, G. (1995), "Eliciting knowledge from experts: a methodological analysis", *Organizational Behavior and Human Decision Processes*, Vol. 62, pp. 129-58.

Holmes, J.H., Durbin, D.R. and Winston, F.K. (2000), "Learning classifier system: an evolutionary computation approach to knowledge discovery in epidemiologic surveillance", *Artificial Intelligence in Medicine*, Vol. 19, pp. 53-74.

Hutchins, E. (1995), *Cognition in the Wild*, MIT Press, Cambridge, MA.

Huzel, D.K. and Huang, D.H. (1992), *Modern Engineering for Design of Liquid-Propellant Rocket Engines*, American Institute of Aeronautics and Astronautics, Washington, DC.

Jurkiewicz, C.L., Knouse, S.B. and Giacalone, R.A. (2001), "When an employee leaves: the effectiveness of clinician exit interviews and surveys", *Clinical Leadership and Management Review*, Vol. 15, pp. 81-4.

Kransdorrf, A. (1997), "Fight organizational memory lapse", *Workforce Magazine*, September, pp. 34-9.

Klein, G. (1992), "Using knowledge engineering to preserve corporate memory", in Hoffman, R.R. (Ed.), *The Psychology of Expertise: Cognitive Research and Empirical AI*, Erlbaum, Mahwah, NJ., pp. 170-90.

Klein. G.A., Calderwood, R. and MacGregor, D. (1989), "Critical decision method for eliciting knowledge", *IEEE Transactions on Systems, Man, and Cybernetics*, Vol. 19, pp. 462-72.

Motta, E., Rajan, Y. and Eisenstadt, M. (1990), "Knowledge acquisition as a process of model refinement", *Knowledge Acquisition Journal*, Vol. 2, pp.21-49.

Novak, J.D. and Gowin, B. (1984), *Learning How to Learn*, Cambridge University Press, New York, NY.

Paris, E. (2001), "Lords of the ring: the fight to build the first US electron-positron collider", *Historical Studies in The Physical and Biological Sciences*, Vol. 31, pp. 355-80.

Regoczei, S.B. and Hirst, G. (1992), "Knowledge and knowledge acquisition in the computational context", in Hoffman, R.R. (Ed.), *The Psychology of Expertise: Cognitive Research and Empirical AI*, Springer Verlag, New York, NY, pp. 12-25.

Seifert, C., Patalano, A., Hammond, K. and Converse, T. (1997), "Experience and expertise: the role of memory in planning for opportunities", in Feltovich, P., Ford, K. and Hoffman, R.R. (Eds), *Expertise in Context*, The MIT Press, Cambridge, MA, pp. 102-23.

Schwartz, C. (1997), "The corporate memory: information management in the electronic age", *Information Processing and Management*, Vol. 33, pp. 810-12.

Smith, K A., Willis, R.J. and Brooks, M. (2000), "Analysis of customer retention and insurance claim patterns using data mining: a case study", *Journal of the Operational Research Society*, Vol. 51, pp. 532-41.

Vicente, K.J. (1999), *Cognitive Work Analysis: Toward Safe, Productive, and Healthy Computer-based Work*, Erlbaum, Mahwah, NJ.

Von Krogh, G., Ichijo, K. and Nonaka, I. (2000), *Enabling Knowledge Creation: How to Unlock the Mystery of Tacit Knowledge and Release the Power of Innovation*, Oxford University Press, New York, NY.

Waterman, D.A. (1986), *A Guide to Expert Systems*, Addison-Wesley, Reading, MA.