Knowledge Modeling and the Creation of El-Tech: A Performance Support and Training System for Electronic Technicians[1]

John. W. Coffey, Alberto J. Cañas, Thomas Reichherzer, Greg Hill,
Niranjan Suri, Roger Carff, Tim Mitrovich & Derek Eberle
Institute for Human and Machine Cognition
University of West Florida
Pensacola FL 32502
www.ihmc.us

Abstract

In this paper, we describe a unique approach to the creation of a system to provide electronic performance support and training for electronics technicians. This work starts with a survey of relevant approaches to knowledge elicitation and modeling for performance support, and a review of other systems that have been created to assist with electronics troubleshooting. We then describe a system named El-Tech (Electronic Technician) that was created as part of a joint research effort with the Chief of Naval Education and Training.

Keywords

knowledge elicitation, knowledge modeling, concept maps, performance support system

[1] Published in: Coffey, J.W., A. J. Cañas, T. Reichherzer, G. Hill, N. Suri, R. Carff, T. Mitrovich & D. Eberle, *Knowledge Modeling and the Creation of El-Tech: A Performance Support and Training System for Electronic Technicians*, Expert Systems with Applications, 25(4) (2003).

1. Introduction

This paper describes a novel approach to the creation of a performance support and training system to assist electronics technicians. The approach to creation of this system involves the collaboration of a domain expert and a knowledge engineer in the development of a multimedia model from which an expert advisory system can be culled (Ford et al., 1996). The knowledge model is retained and used as an explanatory component in conjunction with the inference component of the deployed system (Ford et al. 1991; Ford, Cañas, & Coffey, 1993), as a reference resource to support job performance, and as instructional content for courses on the knowledge domain (Coffey & Canas, 2001).

The system described in this report is named El-Tech. El-Tech was constructed in a joint research effort with the Chief of Naval Education and Training that is aimed at the creation of a knowledge-based performance support system for the Navy's electronics technicians. At any given time, the Navy has approximately 45,000 students in training, mainly through courses taught by approximately 6,500 instructors. It is typically the case that Navy electronics technicians find themselves in-fleet and responsible for mission-critical equipment months after receiving rapid-fire, short duration training and after an extended period of time with no responsibilities for the equipment.

The El-Tech system is designed for use by Naval electronics technicians. The goal of this system is to combine elements of instructional technology with performance support in order to get the right information to the electronics technician at the right time (Wehrenberg, 1989). El-Tech is designed to achieve this goal by assisting electronics technicians in their job performance with two components:

- A browsable knowledge model regarding structures and functions of the equipment, failures, and troubleshooting strategies, and
- an interactive inferencing component that presents the technician with a systematic approach to the diagnosis of equipment failures.

El-Tech was created by using the PreSERVe method (Prepare, Scope, Elicit, Render, Verify) of knowledge modeling (Coffey, Hoffman, Canas & Ford, 2002), an iterative approach that has a high level of acquirability coupled with good expressiveness of the acquired knowledge (Ford & Bradshaw, 1990). This approach to knowledge modeling was developed in work with the National Aeronautical and Space Administration, Glenn Research Center (Coffey, 1999). The knowledge model created with this method was used to create the advisory component system of El-Tech. The knowledge modeling approach described here exploits multiple knowledge acquisition strategies, and opportunistically utilizes information resources that the expert identifies during the process, in order to create an informal but semantically rich model of the knowledge domain.

The rest of this paper presents a brief review of knowledge elicitation and modeling techniques, and a description of strategies and systems that have been developed to assist the electronics technician. This account contains a description of how various knowledge elicitation techniques were chosen and combined in the process of creating an expert knowledge model as the basic component of El-Tech. The process of creating the inference component is also described. A discussion of how the knowledge model created during the knowledge elicitation phase contributes to the capabilities of the deployed system is included.

2. A Survey of Knowledge Elicitation and Knowledge Modeling Techniques

This section contains an overview of methods of knowledge acquisition and basic approaches to knowledge modeling. In the most general sense, approaches to knowledge elicitation can be categorized as direct, in which interactions with one or more domain experts occur, or indirect, in which knowledge is culled from texts, reports, or other documentation (Waterman, 1986). This following section primarily addresses direct approaches. Knowledge modeling approaches start by seeking to provide useful representations of elicited knowledge at the knowledge level - independent of any commitment to a symbolic, machine-usable representation (Moreno et al, 2001). From these representations, automated inference components may be created.

2.1 Knowledge Acquisition Techniques

Many direct knowledge acquisition strategies have been identified including Structured Interviews (Wood et al., 1995), Critical Decision Method (Klein, Calderwood & MacGregor, 1989; Crandell and Getchell-Reiter, 1993), Knowledge Audits (Hoffman, Coffey & Ford, 2000), Cognitive Modeling (Hoffman, Coffey & Ford, 2000), Protocol Analysis (Ericsson & Simon, 1993), and Work Patterns Analysis (Vicente, 1999).

Hoffman, Shadbolt, Burton, & Klein (1995), describe a range of general techniques that include analysis of the tasks experts perform as they engage in problem solving or decision-making, the use of unstructured or structured interview techniques, and the employment of so-called "contrived" techniques. The following brief discussion, taken from Hoffman et al. (1995), describes these general categories.

The analysis of familiar tasks encompasses a range of activities that involve observation of the expert actually performing work. In this approach, the expert typically verbalizes while performing a task. This procedure can lead to the enumeration of a protocol for the performance of the task. This method utilizes test cases that may be typical of the sort of problems the expert faces, or that may be anomalous in the sense that they occur infrequently, or because they are difficult to solve.

Unstructured interviews are informal conversations geared toward a broad goal. They can be useful early in the process when basic, groundwork-laying information is sought.

However, some degree of structure is typically best for all but initial interviews. Interviews may have varying degrees of structure along at least two dimensions: a) the type of question asked and b) adherence to a pre-defined set of questions. Types of questions can include (GAO, 1991):

- Open-ended - essentially having no constraint on the answer
- Fill-in-the-blank - require a single or multi-word answer
- Binary - yes-no or true-false
- Scaled-response - utilize Likkert scale type questions

The degree of adherence to predefined questions can range from asking only the predetermined questions in exactly the predefined order, to the utilization of a subset of the original questions as starting points for subsequently unstructured discourse. Structuring questions keep the discussion from meandering, but may fail to address important issues that are not on the agenda.

Contrived techniques such as decision analysis, rating and sorting tasks, graph construction, and constrained processing/limited information problems require the expert to perform other activities than their familiar tasks. 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 determined and represented in a reasoning model. Rating and sorting tasks involve making judgments regarding attributes of a problem domain.

Graph construction involves the creation of conceptual graphs (Sowa, 1992) or concept maps (Novak and Gowin, 1984) that are structured, non-textual representations of knowledge based upon Ausubel's (1968) Assimilation Theory. Concept maps are graphs that are comprised of concepts on 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 & Bradshaw, 1995) and providing a framework for the structuring of knowledge (Cañas, Ford & Coffey, 1994).

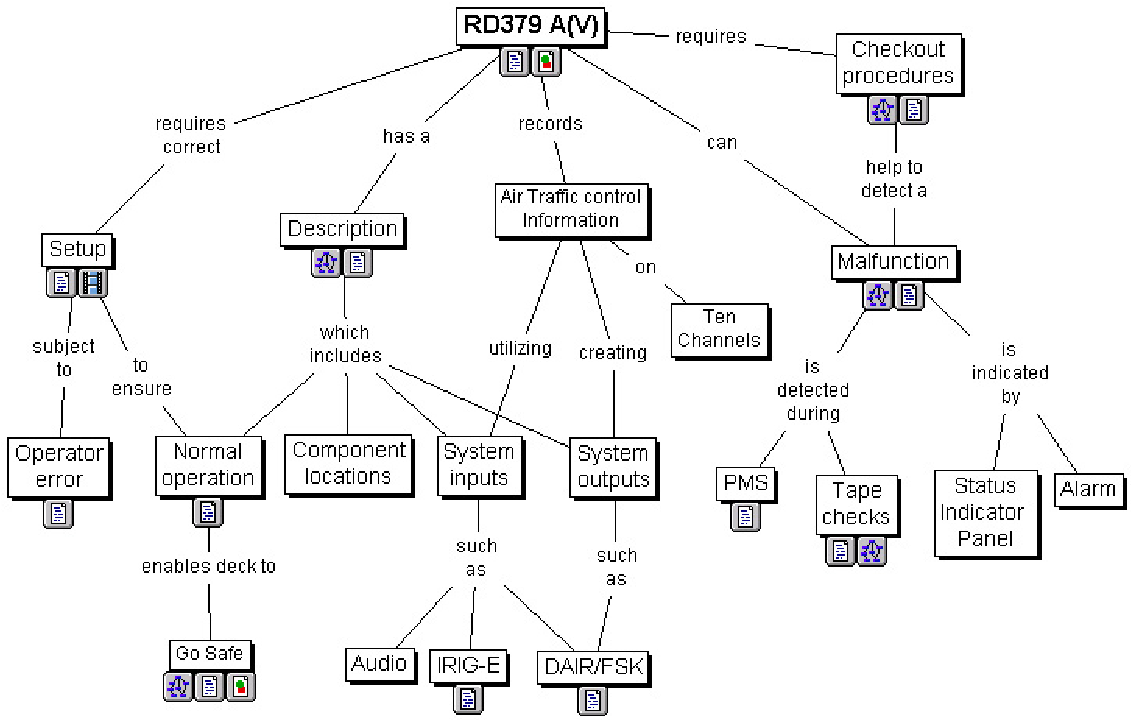

Figure 1. The most general concept map from El-Tech

Figure 1 presents the most general concept map from El-Tech. The concepts appear in the shadowed boxes and the linking phrases, representing the relationships among concepts, are the labels on the arcs. Concept maps are especially useful for the determination of the scope of a knowledge elicitation effort by elaborating the most basic concepts to be considered. Concept maps also support Aussubel's (1968) notion of subsumption in which more detailed concepts are subsumed under more general ones. Interview structure can be based on the elicited general, subsuming concepts in the concept maps. The act of working through the issues in the context of the concepts from Concept Maps creates an efficient framework for structured discussions. Note that the Concept Map in Figure 1 is augmented with icons attached to many of the concepts in the map. These represent additional resources such as texts, graphics and video that are pertinent to the concept and part of the knowledge modeling process described in more detail in section 3.4.

The three broad categories of knowledge elicitation techniques described here can be used in various combinations and with a focus on eliciting knowledge of a certain type in order to serve a specific purpose. For example, less structured interviews can be used initially to lay groundwork such as the identification of the expert or experts and the knowledge domain. Concept mapping helps to define the initial boundaries around the effort. Structured interviews and analysis of familiar tasks can be used to articulate details, and retrospection on memorable cases can help to identify anomalous or difficult cases.

2.2 Knowledge Modeling Approaches

This section contains a survey of representative knowledge modeling approaches. Knowledge modeling methods and languages may be thought of as representation schemes that augment traditional data modeling by adding semantic content to the modeling language (Mineau et al. 2000). Knowledge modeling approaches lie on a continuum from informal (and potentially, easily understood by humans) to formal (and therefore, capable of being evaluated by machine). Chan and Johnston (1996) describe two categories of approaches to knowledge modeling: one group based upon problem solving methods and another based upon domain ontologies. These two approaches have significant overlap in the sense that, although problem solving methods are process oriented and ontological accounts start with characterizations of objects, at some point during knowledge model construction, process models must be created.

The KADS methodology places emphasis on the structure of tasks needed to accomplish a goal, and is typical of what Chan and Johnston would characterize as a problem solving method. A KADS system is represented as a group of models, each of which represents some part of the system (Schreiber, Akkermans & Hoog, 1994). The CommonKADS methodology (Schreiber et al, 2000) contains a six-level, hierarchical representation of: the overall domain, large-scale processes, agents and communications involved in the processes, the expertise to perform the processes, and design of the system. By adding the notion of domain knowledge to process views of the domain, CommonKADS effectively creates an ontological account of the knowledge domain.

Moreno et al. (2001) describe a case study in the use of KADS for the creation of a hospital management system. Moreno et al describe the two major components of the knowledge model: domain knowledge and control knowledge. Domain knowledge includes items such as services the hospital provides, human and material resources the hospital has, specialties, patients, tests, etc. Control knowledge pertains to the reasoning processes of the system. Moreno et al. started the knowledge modeling process by creating a conceptual graph of the hospital system. A conceptual graph is like a concept map with a constrained set of linking phrases such as "a_part_of" and "a_kind_of." From this conceptual graph, they defined a Frame-ontology for their hospital management system. The idea of an ontology in the computer/cognitive science sense is a precise specification of a common parlance or vernacular for a knowledge domain (Gruber, 1993). Ontolingua is a tool that supports the creation of domain-specific ontologies that can be specified in a formal way.

2.3 Knowledge-Based Systems to Aid with Electronics Fault Diagnosis

Rafea, El-Desouki and El-Moniem (1990) describe two basic approaches to building diagnostic expert systems The shallow knowledge or “fault model” approach and the structural-functional or “deep model” approach. With the fault model approach, knowledge engineers attempt to capture the heuristic knowledge (rules of thumb) of the expert. Such systems are typically domain-dependant but efficient, both in the speed with which they can be constructed and the speed with which they can converge on a failure when consulted. With the deep or “structural-functional” model approach KEs seek to create a generic model of nominal performance and failures that are more general. Rafea, El-Desouki and El-Moniem describe benefits and drawbacks of each approach, and a system that is a hybrid of both approaches.

Clancy (1987) describes a qualitative reasoning approach to the detection of faults in switch mode power supplies (SMPSs). This approach attempts to capture the usual way that a technician would decide on the cause of the failure, by looking for basic presence or absence of signals. The system described by Clancy simplifies the diagnostic process by not requiring the detailed analysis of signals. This approach exploits similarities in SMPSs to create a general approach to fault diagnosis in these components.

A system described by Cunningham (1998) contains a generic deep model to deal with SMPSs. The system contains structural and behavioral knowledge about modules and components of such circuits. These so-called “building blocks” can be recombined to describe specific power supply designs in a way that leverages the generic elements. The system reported was able to locate between 80% and 90% of faults in these power supplies.

Renfrew and Tian (1993) also describe a specifically targeted system that deals with three-phase electronic inverters. These circuits have similar switch configurations and for that reason, can be tested in more generic ways than other circuits. The authors tested circuits of this type with faults seeded in the systems to find differences in outputs. From the results of the tests, the authors were able to create a rule-based system that could generically diagnose such circuits from outputs.

Redford (1992) presented work in the area of analog electronics creating a searchable, hierarchical classification system based upon manufacturers’ data sheets. A variety of search methods are employed to reason through faults. Strategies include hill-climbing algorithms, problem reduction and constraint satisfaction. The author makes the claim that the described system is an implementation of the way that human electronics technicians reason.

The systems described here have all had the purpose of trying to create systems to assist with diagnosis of electronic circuit failure. As the literature shows, attempts have been made to create systems that are based upon heuristics known to expert diagnosticians, upon deep models of structural knowledge, and upon attempts to leverage specific attributes of specific types of circuits. As would be expected, it appears that no one approach clearly proves to be ideal for all circuits and situations.

3. The El-Tech System

The El-Tech system models the knowledge of an expert electronics technician on the RD-379A(V)/UNH, a fault-tolerant air traffic control recorder/reproducer manufactured by

Magnasync/Moviola corporation. This equipment tracks all flight data for aircraft in the vicinity of an airport or aircraft carrier. It contains a real-time clock as well as many other analog and digital electronic components. This section contains a description of the use of the method used to create a knowledge model of the RD-379A(V), the method of creating the inferencing component from the model, and a description of the deployed system.

3.1. Creating El-Tech

The basic approach to the creation of El-Tech was to create an expert knowledge model and to create the inference component from the model. The knowledge modeling effort led to the creation of an integrated rendering of conceptual and process knowledge of the domain that defined the boundaries of the system, and the creation and verification of the consultation component of the system from knowledge contained in the model. The PreSERVe method (Prepare, Scope, Elicit, Render, Verify), an iterative method of knowledge modeling (Coffey, Hoffman, Canas & Ford, 2002) was employed in the creation of the knowledge model. The next sections describe how this method was applied in the overall process, and how the consultation component was culled from the knowledge model that was constructed.

3.1.1 Knowledge Model Creation

Preparation for the creation of El-Tech had several dimensions. It was necessary to choose an expert who would be available, obviously precluding sailors who were in-fleet. It was decided to utilize a chief petty officer from the training school who had extensive in-fleet experience. Most of the early interactions with the Navy representatives who were involved in the project were based upon unstructured interviews. On the basis of such interviews, the system of interest was chosen from among several possibilities. In another part of the preparation phase, Knowledge Engineers (KEs) reviewed basic electronics and sat in on classes pertaining to the system of interest, in order to improve their background knowledge of the domain. The time frame for the project was short, so the initial scope of the project was kept relatively narrow. The scope of the project broadened as the project went forward.

Knowledge elicitation commenced with a contrived technique, the creation of concept maps for the domain. Figure 1 contains a depiction of the final top-level concept map of the system. This concept map had gone through several iterations during the process. The KEs originally thought the system would be strictly focused on troubleshooting, and it was toward the creation of a troubleshooting model that knowledge elicitation began. In the process, the expert realized and revealed that system setup and checkout procedures often lead to the identification of problems and should be included in the knowledge model. These ideas were made explicit in the model, as can be seen in the map in Figure 1. The importance of these other ideas caused a broadening of the scope of the project to encompass setup and checkout procedures. The need to be vigilant in the ongoing assessment of the scope of a knowledge modeling endeavor was made quite clear by the broadening of the scope to include the three dimensions (setup, checkout and malfunctions).

Another broadening of the scope occurred since it was deemed desirable to create a basic description of the RD-379A(V)/UNH that documents its components, their locations, and how they individually and collectively operate. In conjunction with this effort, the setup, checkout, and maintenance procedures were identified and described in the system. The third area of knowledge elicitation involved identifying the symptoms of problems that might be noted by the technician, and the diagnostic strategy that might be employed to repair the problem.

After the three areas were identified, knowledge elicitation proceeded with the creation of concept maps to represent the important concerns of the expert in each area. An iterative process of refining the concept maps ensued. Model rendering proceeded with the refinement of the concept maps and with creation of support materials such as schematics, block diagrams, video of the expert discussing setup and troubleshooting techniques, etc. Another important part of model rendering was the establishment of navigational links among the various components (map to map and map to resource) of the emerging knowledge model.

One aspect of the resource-rendering component of this work was relatively simple since technical manuals were available. However, the manuals were in hardcopy form and they had to be converted to digital form for inclusion. It is typical that significant resources exist in some sort of hardcopy format. The effort to convert such resources to digital format is often defensible since the resources gain a lot of added utility both for performance support and for training. Other decisions were made regarding the best uses of resources that would be created from scratch. Literal graphics were captured as photographs, schematics from books were digitized, etc. Of special concern was the creation of digital video. Creating and editing digital video is costly, as is the bandwidth to stream it effectively over a network. The governing principle was to create videos for issues that could not be explained adequately by other media.

Verification of the emerging model was also iterative. A malfunctions map helped to identify the major areas of malfunctions including heuristic knowledge of the more or less commonly occurring ones. Verification started with review of the concept maps themselves and proceeded to accompanying resources and links between resources.

Table 1. Summary of knowledge elicitation techniques utilized in creating El-Tech.

Unstructured Interview	Structured Interview	Contrived Techniques	Analysis of Familiar Tasks	Retrospections
Little (initially)	High	High	Moderate-High	Moderate

Table 1 presents a summary of the degree of reliance on the various knowledge elicitation techniques that were used in the creation of El-Tech. The use of the contrived technique of creating concept maps to assess the scope of the project and to lend structure to interviews was critical. The knowledge model also served as a benchmark to help assess that a comprehensive evaluation of pertinent cases (analysis of familiar tasks) had been carried out. Use of retrospection to consider difficult or anomalous cases akin to the critical decision method, was an important, but not extensively used method.

3.1.2. Building the Inferencing Component.

The inference component of El-Tech is a rule-based advisory system, with rules culled from the knowledge model elicited from the expert. The Fault Isolation knowledge bases were first developed in CLIPS (Riley, 1997). CLIPS rules could easily be translated to JESS (Java Expert System Shell) (Friedman-Hill, 1997), which allowed creation and maintenance of centralized knowledge bases on a network, and a graphical user interface that provides the front end to the knowledge model. A distributed collection of servers house these components (the knowledge model and the inference component) together with all the other media, making the entire system available anywhere in-fleet via a network connection. An alternative would be to deploy the system on CD.

The creation of concept maps and accompanying resources led to the elicitation of comprehensive conceptual and process knowledge of the domain. Although the goal of eliciting concept maps is to create a conceptual rather than a procedural model, a concise rendering of the problem solving processes of the expert is explicit in the concept maps.

Both very fundamental distinctions and highly detailed heuristics that were pertinent to the inferencing component were identified as a result of creating the concept maps. A very basic heuristic was the fact that problems can be detected either from an alarm triggered by the system itself, or in the course of performing the checkout procedures. The system had no on-board diagnostics other than the ability to sound an alarm if something went wrong. The second very basic item of interest that was identified during the map creation process was the issue of whether the failure had occurred on one or both of the machine’s transports. This heuristic was the most basic of many heuristics that allowed for rapid fault isolation.

The emerging conceptual model was augmented by the analysis of familiar tasks, in this case the analysis of typical scenarios of problem symptoms. Descriptions of such scenarios were attached as accompanying resources to nodes in the concept maps. Retrospection relative to anomalous cases was employed in order to elicit information regarding less commonly seen problems.

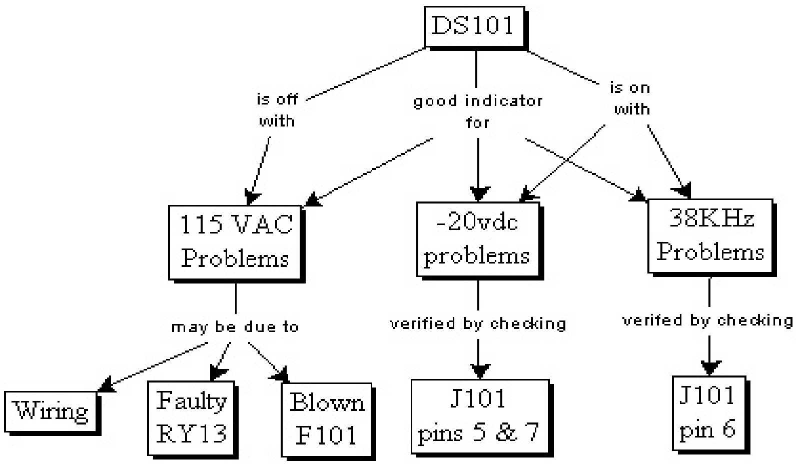

Figure 2. A map fragment on troubleshooting the Bias Oscillator in El-Tech

In the process of map creation, highly detailed heuristics were identified as well. Figure 2 contains a fragment of a concept map (without the navigation icons) on the Bias Oscillator of the machine. Although this map contains essentially conceptual rather than procedural information, it displays two very useful aspects:

- It conveys the rapid, shortcut causal rules (heuristic knowledge) that differentiates expert performance from means-ends analysis typical of novice performance, and
- it is highly expressive and easily transformed into rules.

As Figure 2 illustrates, the propositional nature of concept maps make explicit many of the underlying rules for problem solving. This concept map is describing a way to use the power light on the bias oscillator in conjunction with tests on a few pins of component J101 to detect very quickly whether a power problem is in a 38KHz circuit or in a –20vdc circuit. Making the determination that a power problem in the bias oscillator was a 38KHz or a –20vdc problem would potentially take a long time if a relatively novice sailor or an automated inferencing system attempted to reason from first principles or with a deep knowledge model of electronics and a schematic diagram. Figure 3 presents a rule that was derived from the concept map in Figure 2.

```
(defrule BiasOscillator_J101_38kHzCheck
    (goal
      (type J101_pins5and7)
      (value "BothGood"))
  =>
    (printout t "Be sure to switch to AC for 38KHz check." crlf)
    (printout t "Does pin 6 have 38KHz, 40vpp on J101 (y/n)?" crlf)
    (assert (goal
      (type J101_pin6)
      (value (readline))))
 )
```

Figure 3. A CLIPS/JESS rule developed from the Concept Map in Figure 4.

3.2 The Deployed System

The deployed system has two basic components, an interactive question-answer type consultation component and the multi-media knowledge model to explain the consultation and serve as instructional content. A knowledge model that is comprehensive and can account for all the knowledge needed to create the inferencing component of a performance support system also contains all the conceptual knowledge to explain the consultation component's questions and actions. (Ford & Bradshaw, 1995).

Figure 4 presents a graphic of the system interface. The leftmost window in Figure 4 depicts an interaction with the advisory component of the system, in which the system is asking the user about the machine's performance in "direct mode." The system explanation of the question has been invoked and the "direct mode", as it pertains to a malfunction, has been located in the "Malfunctions" concept map. The user has opened a window that gives a textual description of the importance of direct mode in the diagnosis of the fault, and is currently viewing a digital video of the expert elaborating on the point.

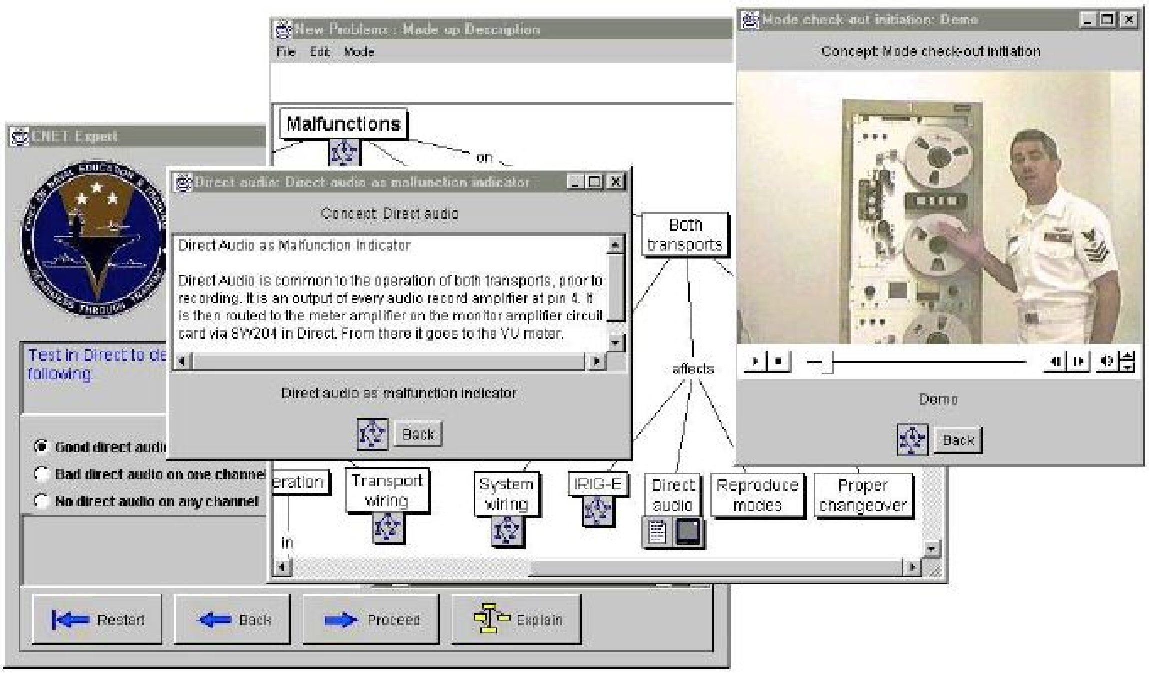

Figure 4. The Interface of the deployed system.

The malfunction model is closely coupled with a model that describes the nominal operation of the system so that the user can easily navigate through a description of a component's normal behaviors and failure modes. The knowledge model contains two hierarchies of concept maps, the first presenting the detailed description of the equipment, and the other the detailed account of how to diagnose and repair faults. These two constituents of the knowledge model are tied together by the top concept map depicted in Figure 1. The knowledge model includes block diagrams of the equipment, electronics schematics, photos that show components and their locations, textual passages transcribed from interviews with the expert, digital video that illustrates checkout and fault isolation procedures, and references to supplementary information contained in external sources.

Users can utilize the system in either of two ways: by invoking the advisory system to assist in the diagnosis of problems, or as an instructional system. When the user requests an advisory consultation, the system asks a series of questions concerning symptoms to be found in the equipment and the knowledge model serves an explanatory function to explain why the system is asking the user a question. If the user utilizes the knowledge model for training, instruction or review, the knowledge model serves as a browsable, hyperlinked information resource.

When the user initiates a consultation, the system asks questions regarding the symptoms of the problem and adjusts its line of reasoning based upon the answers it receives. Along the way, the inferencing component suggests tests the user should make on the equipment, and makes explicit safety recommendations (for instance, the need to use heat sinks when testing, etc). In addition, the system provides guidance on carrying out the test. As an example (as can be seen in Figure 3), the system reminds the electronics technician to switch from dc to ac settings on the test equipment before making the test. The expert knew that this was a common mistake made by novice technicians and the goal of enhancing performance would dictate that this suggestion be made explicit.

The inference component converges on probable failure points at the component level, individual transistors, relays, switches, motors, amplifiers, or even a wiring fault, typically after the user has answered 3 to 7 questions. At any time during the consultation, the user can ask for an explanation of the line of reasoning the system is pursuing. Upon such a request, the inference component transports the user into the knowledge model to a place that contains media that relate to the question. Such media might include a concept map that pertains to the issue, text descriptions of the knowledge behind the question, video explanations by the expert, graphics of the schematics and literal pictures that show test points of the components under consideration, oscilloscope settings, wave forms, test point values, etc.

The knowledge model can also be used as a freely browsable hypermedia that provides a description of the basic system as it operates nominally, a description of typical and anomalous malfunctions, and links between these knowledge model components. The system is based upon actual, real-world experiences of the expert rather than on engineering reports that are decoupled from operational problems with the equipment. The knowledge model is of use both in the school house during initial training, and as a performance support system for in-fleet sailors.

4. Summary and Discussion

4.1. Summary

This work contains a presentation of an overview of knowledge elicitation methods as they can be employed in the production of a model of a knowledge domain. The type of knowledge model described here is based upon concept maps and, although it is less formal than some representation schemes, it is semantically very rich and easily understood by non-technical people.

This work has resulted in production of a prototype system (El-Tech) that assists electronics technicians in their training and job performance. The development and deployment of systems like El-Tech will make it possible for electronics technicians to have access to expert, field-tested knowledge when and where it is needed, both during initial training and in-fleet. This knowledge modeling approach is general and is currently being used in another project with the National Imagery and Mapping Agency (NIMA).

This research has led to the refinement of a knowledge modeling method named PreSERVe (Coffey, et al, 2002), a method of initial preparation, followed by an iterative process of examining scope of the project, eliciting knowledge, rendering the knowledge into a hypermedia knowledge model, and verifying the emerging model.

4.2. Discussion

4.2.1 Knowledge Elicitation for Knowledge Modeling

Although the entire gamut of knowledge acquisition schemes are compatible with a knowledge modeling approach to knowledge elicitation, some methods have greater utility than others in specific situations and at different phases of the effort. For instance, unstructured interviews had utility initially, but declined in importance as knowledge elicitation progressed. Concept mapping is a highly efficient means of ascertaining and making explicit the scope of the effort, and serve as a minimally biasing basis for interviews with varying degrees of structure. Analysis of familiar tasks is helpful with creation of process models, and retrospection on memorable cases helps to identify the more difficult or anomalous cases that are often the most difficult for a journeyman practitioner to resolve successfully.

The iterative nature of the knowledge modeling method that was employed proved to be an important aspect. Additionally, the emphasis on reassessing the scope of the project was deemed critical in knowledge elicitation where much of the most important knowledge of the expert is tacit and only uncovered as the effort goes forward. The creation of concept maps leads to elaboration of heuristic knowledge that casts light on the efficient reasoning strategies employed by the expert. This paper presented several examples of the power of such heuristic knowledge.

Knowledge modeling of this sort integrates conceptual and process knowledge. Such an integration is beneficial at two levels. At a functional level, it permits the system to direct the user to the conceptual knowledge upon which the inference process is based as the user holds a consultation with the system. More generally, if the user is browsing through the knowledge model to learn about the system, the model’s training value is enhanced by having conceptual knowledge clearly associated with process knowledge.

4.2.2 Formal vs informal knowledge modeling

As described in the literature review, this knowledge modeling approach contrasts with others based upon more formal knowledge representations such as conceptual graphs or ontolingua representations. This representation trades off being highly expressive in machine terms for the ensuing gains in comprehensibility in human terms. Such gains make the explanatory and instructional capabilities of the system much more rich, and foster the reuse of content based upon expert knowledge as training materials.

The tradeoff is that systems such as this are quite labor-intensive to make and maintain. The knowledge acquisition bottleneck is by now, a well-known and seemingly intractable difficulty against which one can only expect to make incremental gains. However, experience has convinced the authors that concept mapping as a central component of a principled knowledge modeling method such as PreSERVe can clearly provide incremental gains in the process.

4.2.3 Heuristic vs Structural-Functional models

The El-Tech system illustrates an approach to performance support systems that combines many of the best elements of both heuristic and structural-functional models of a problem domain. The part of the knowledge model that pertains to the nominal performance of the system is essentially a structural-functional model. Malfunction modes are clearly elaborated. Augmenting this part of the system is the inferencing component that converges so rapidly because of the heuristic knowledge it employs.

Since the structural-functional model forms the explanation component of the (heuristic-based) consultation component, the user is afforded elements of both approaches. In terms of instruction or training, the freely browsable knowledge model combines elements of a structural-functional model, essentially the first principles of the domain, intermingled with heuristics pertaining to the model. Such a model allows the user to benefit from the yeas of experience the expert has accumulated in the domain.

4.2.4 Specifically Targeted versus Generalizable Systems

Several of the diagnostic systems for electronics that were described in the literature review sought to exploit specific characteristics of specific types of electronic equipment. Quite clearly, these are one-off, special purpose systems. Although El-Tech contains large amounts of system-specific information, the approach itself is completely generalizable and could be applied to any electronic system. Furthermore, while content pertinent to basic electronics was only minimally incorporated into El-Tech, it could be incorporated once and then reused in any subsequent systems that are built.

4.2.5 Final Thoughts

Just as various knowledge acquisition methods are compatible with the knowledge modeling method employed here, other forms of reasoning are compatible as well. Future work will explore the feasibility of incorporating a case-based reasoner into El-Tech. System problems in the Navy trigger an OpNav47902K trouble report that details the nature of the problem, how and when it was discovered, length of time to repair, etc. The Navy has an extremely comprehensive record of these reports. A retrieval system for previous failures could prove a useful adjunct to the current system. A set of retrieved cases would serve the function of an accompanying resource associated with a concept in a concept map, similar to the texts, graphics and video depicted in Figure 4. It is anticipated that application of this technology to other knowledge domains might well reveal other domain-specific ancillary services that can be integrated into this type of system.

References

Ausubel, D.P. (1968). *Educational psychology: a cognitive view*. New York, NY: Rinehart and Winston.

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

Clancy, C. (1987). Qualitative reasoning in electronic fault diagnosis. *Electronic Engineering* 59(731) 141-142.

Coffey, J.W. (1999). Knowledge Preservation at Glenn Research Center. *Unpublished Technical Report,* Glenn Research Center, Cleveland, OH.

Coffey, J.W., & Cañas, A.J. (2000). A Learning Environment Organizer for AsynchronousDistance Learning Systems. *Proceedings of the Twelth IASTED InternationalConference Parallel and Distributed Computing and Systems (PDCS 2000).* November 6-9, 2000 Las Vegas, Nevada.

Coffey, J.W., Hoffman, R.R., Canas, A.J., & Ford, K.M. (2002). A Concept Map-Based Knowledge Modeling Aapproach to Expert Knowledge Sharing. *Proceedings of IKS2002, IASTED International Conference on Information and Knowledge Sharing*. St. Johns Virgin Islands, Nov 6-8, 2002.

Crandell, B., & Getchell-Reiter, K. (1993). Critical decision method: A technique for eliciting concrete assessment indicators from the "intuition" of NICU nurses. *Advances in Nursing Sciences*, 16(1), 42-51.

Cunningham, P. (1998). A case study on the use of model-based systems for electronic fault diagnosis. *Artificial Intelligence in Engineering*. 12(3), 283-295.

Ericsson, K.A., & Simon, H.A. (1993). *Protocol analysis: Verbal reports as data*. Cambridge, MA: The MIT Press.

Ford, K.M., & Bradshaw, J. (1990). Knowledge Acquisition for Expert Systems. *A Workshop presented at FLAIRS-90, The Third Florida AI Research Symposium*, Cocoa Beach, FL April 2-5.

Ford, K.M., & Bradshaw, J. (1995). Beyond the repertory grid: New approaches to constructivist knowledge acquisition tool development. *International Journal of Intelligent Systems*, 8(1), 287-333.

Ford, K.M., Cañas, A.J., & Coffey, J.W. (1993). Participatory Explanation. *Proceedings of 6th Annual Florida AI Research Symposium* (pp. 85-90). Ft. Lauderdale, FL.

Ford, K.M., Coffey, J.W., Cañas, A.J., Andrews, E.J., & Turner, C.W. (1996). Diagnosis and Explanation by a Nuclear Cardiology Expert System. *International Journal of Expert Systems*, 9(4), 499-506.

Ford, K.M., Stahl, H., Adams-Webber, J.R., Cañas, A.J., Novak, J.D., & Jones, J.C. (1991). ICONKAT: An integrated constructivist knowledge acquisition tool. *Knowledge Acquisition Journal*, 3, 215-236.

Friedman-Hill, E. (1997). Jess, the Java Expert System Shell. [online reference] http://herzberg.ca.sandia.gov/jess/.

General Accounting Office. (1991). Using Structured Interviewing Techniques. Washington, D.C.: Program Evaluation and Methodology Division (http://www.gao.gov/policy/10_1_5.pdf).

Gruber, T. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2), 199-220.

Hoffman, R.R., Coffey, J.W. & Ford, K.M. (2000). A Case Study in the Research Paradigm of Human-Centered Computing: Local expertise in weather forecasting. *Unpublished Technical Report, National Imagery and Mapping Agency*. Washington, D. C.

Hoffman, R.R., Shadbolt, N.R., Burton, A.M., & Klein, G. (1995). Eliciting knowledge from experts: A methodological analysis. *Organizational Behavior and Human Decision Processes*, 62(2), 129-158.

Klein. G.A., Calderwood, R., & MacGregor, D. (1989). Critical Decision Method for eliciting knowledge. *IEEE Transactions on Systems, Man, and Cybernetics*, 19(3), 462-472.

Mineau, G.W., Missaoui, R., & Godinx, R. (2000). Conceptual modeling for data and knowledge management. *Data and Knowledge Engineering*, 33, 137-168.

Moreno, L. Aquilar, R.M., Pineiro, J.D., Estevez, J.I., Sigut, J.F., & Gonzalez, C. (2001). Using KADS methodology in a simulation assisted knowledge based system: application to hospital management. *Expert Systems with Applications*, 20, 235-249.

Novak, J.D. & Gowin, D.B. (1984). *Learning How to Learn*. Ithaca, NY: Cornell University Press.

Redford, M.A. (1992). Techniques in electronic diagnosis. Proceedings of DEXA '92, *The International Conference on Database and Expert Systems Applications*, 2-4 Spetember, Valencia, Spain.

Rafea, A., El-Desouki, A., & Abd El-Moniem, S. (1990). Combined model expert system for electronics fault diagnosis. *Proceedings of the Third International Conference on Industrial and Engineering Applications of Artificial Inteligence and Expert Systems*. 1, 32-40.

Renfrew, A.C., & Tian, J.X. (1993). Use of a knowledge-based system in power electronic circuit fault diagnosis. *Proceedings of the 5th European Conference on Power Electronics and Applications*, Brighton, UK. Published by Michael Faraday House, Stevenage, England, 7(377) 57-62.

Riley, G. (1997). CLIPS, A Tool for Building Expert Systems. [online reference] http://www.ghg.net/clips/CLIPS.html.

Schreiber, G., Akkermans, B., & Hoog, R. (1994). CommonKads: a comprehensive methodology for KNBS development. *IEEE Expert*, 9(6), 28-37.

Sowa, J.F. (1992). Conceptual Graphs Summary, in Conceptual Structures: Current Research and Practice, P. Eklund, T. Nagle, J. Nagle, and L. Gerholz, eds., Ellis Horwood, pp. 3-52.

Vicente, K.J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work*. Mahweh, N.J: Lawrence Erlbaum Associates.

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

Wehrenberg, S.B. (1989). The Future Just-In-Time Work Force, *The Personnel Journal*, 36-44.

Wood, L.E., Davis, T.C., Clay, S.L., Ford, J.M., & Lammersen, S. (1995). Evaluation of interviewing methods and mediating representations for knowledge acquisition. *International Journal of Expert Systems Research and Applications*, 8(1), 1-23.