N89 - 15554

The Elements of Design Knowledge Capture

Dr. Michael S. Freeman
National Aeronautics and Space Administration
Systems Engineering Division
Systems Analysis and Integration Laboratory
Marshall Space Flight Center AL 35812

Abstract

This paper will present the basic constituents of a design knowledge capture effort. This will include a discussion of the types of knowledge to be captured in such an effort and the difference between design knowledge capture and more traditional knowledge base construction. These differences include both knowledge base structure and knowledge acquisition approach. The motivation for establishing a design knowledge capture effort as an integral part of major NASA programs will be outlined, along with the current NASA position on that subject. Finally the approach taken in design knowledge capture for Space Station will be contrasted with that used in the HSTDEK project.

Introduction

Although Design Knowledge Capture (DKC) has become a subject of strong interest in the Artificial Intelligence (AI) community, and is a critical element of some major new NASA programs such as Space Station Freedom, there still exists considerable confusion as to its nature, scope and feasibility. This is aggravated by an ambiguity inherent in the term itself, and the natural inclination to view DKC as just a "scaled-up" knowledge based system development task. In this paper I will try to clarify some of the differences between DKC and typical knowledge engineering projects, identify some of the challenges inherent in these differences, describe the approaches being taken in NASA to meet these challenges, and discuss the needs/benefits which justify the effort being made to develop a methodology for DKC. As will be discussed in more detail below, it is my opinion that an ability to integrate DKC with the traditional engineering activities which comprise a NASA system development project will be essential to the success of the missions we are undertaking in the next decade, and beyond. It is also an outstanding opportunity for NASA to develop new technology which can have a pervasive and significant impact on American industry, fundamentally changing the way we perceive technological progress. NASA is responding to this challenge by setting itself the goal becoming a world leader in the area of intelligent autonomous systems for aerospace applications, and has set up an aggressive program for research, development, validation, and demonstration of DKC technology. The Hubble Space Telescope Design/Engineering Knowledgebase (HSTDEK) Project is the primary focus for DKC development in NASA at the present time, although other significant activities exist. As the lead center for the HSTDEK Project, Marshall Space Flight Center has made a strong commitment to the utilization of knowledge based systems in future missions.

The Nature of Design Knowledge Capture

In order to distinguish DKC as a special type of knowledge engineering activity, it will benefit us to first review some basic nomenclature. Artificial Intelligence is "the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit the the characteristics we associate with intelligence in human behavior - understanding language, learning, reasoning, solving problems, and so on". [1] It is in the area of problem solving that AI has seen its greatest practical successes, in the form of "Expert Systems" which emulate in software and hardware the ability of some person (or persons) recognized as having exceptional skills in solving a particular type of problem. The ability of a computer system to mimic such skills can be very valuable, as has been widely discussed in AI, and business, literature of the past few years. The process of capturing such skills in a form ("Knowledge Representations") amenable to implementation in a computer system (a "Knowledge Base") is called "Knowledge Acquisition", and a number of highly sophisticated tools have been developed to support this type of activity. These range from symbolic languages such as LISP and PROLOG, to single-paradigm tools (or "shells") such as Rulemaster and VP-Expert, and on to multi-paradigm development environments such as KEE and ART. The ability of these tools to support development of expert systems and the commercial success of the systems built using those tools have reinforced each other, resulting in better tools for constructing expert systems. But it also has caused the development of these tools to focus closely on the needs of expert systems at the expense of more generalized capabilities. This has, in turn, constrained the pursuit of intelligent systems with broader objectives [2]. These broader focus intelligent systems are called "Knowledge Based Systems" (KBS), of which expert systems are a subset. Where expert systems attempt to capture the scarce expertise available only from one or a very few widely recognized experts, knowledge based systems attempt to formalize the more general problem solving techniques more typical of competent engineers. For this reason they are less tied to a narrow application area, or "domain".

The goal of a DKC effort, as stated for NASA's Hubble Space Telescope Design/Engineering Knowledgebase Project, is to enable major projects to capture the design/engineering expertise they have acquired during the development of their systems in a knowledge base capable of supporting multiple applications [3]. Because of the nature of the expertise to be captured, these applications are knowledge based systems rather than expert systems. The knowledge utilized in developing a major aerospace system spans many technical disciplines, applied by a large number of engineers. In the past, some isolated areas of this expertise might be singled out as especially valuable or unique, and an expert system development project initiated. The premise behind DKC is that, given the complexity, cost, and operational lifetime of current major systems, a much broader and better integrated cross section of the expertise will be required to insure long term mission success.

It is not possible at this point in the development of a DKC methodology to rigorously delimit the scope or type of expertise which must be captured. In order to approach the question at all, we must distinguish between a knowledge base and a knowledge based system. A knowledge base typically consists of a "domain model" of some activity or component of a system coupled with expertise on how the elements of that model interact in some context. A knowledge based system utilizes the contents of the knowledge base, augmented with application specific expertise, to achieve some goal. In expert systems the domain model is so specifically tailored to a specific narrow application that the domain knowledge is difficult or impossible to separate from the application. In other words, the contents of the knowledge base cannot easily be reused to support multiple applications. Since DKC is intended to result in a knowledge base capable of supporting multiple applications, this is a significant difference. One consequence of this difference is that tools developed to build expert systems provide limited support to a DKC effort. Another effect is that a thorough DKC effort can be expected to result in a knowledge base much larger than other knowledge engineering activities.

In addition to inherent large-scale, multi-application characteristics, a DKC knowledge base also must involve the integration of multiple technical disciplines. A typical expert system will focus on only one technical discipline, such as biochemistry, VLSI electronics, or mathematics. DKC, therefore, poses formidable systems engineering (as well as knowledge engineering) problems. A typical NASA development team may represent thirty or more major technical disciplines, with consultants in narrower areas. Between contractor and civil service engineers, there may be hundreds of people involved in the development process. Current technology is adequate to build isolated expert systems in most of these disciplines. The challenge for DKC is to develop a methodology which will integrate the expertise of these various disciplines into a knowledge base in which they can be jointly applied to solve problems which cross discipline boundaries. At present, a systems engineering approach can be facilitated in a DKC project by establishing from the beginning some target applications to be supported by the DKC knowledge base which involve multiple technical disciplines, and address the system at an integrated (as well as component) level. Two good candidates which currently require a great deal of manpower to support in NASA projects are fault diagnosis (including Failure Modes Effects Analysis and Critical Items Lists) and training for operational support. By insuring that the DKC knowledge base can support these specific functions, as well as meet its general knowledge capture goals, we can provide both a near term return on investment and enhanced long term support. The state of the art in knowledge engineering and/or project requirements may result in the choice of other applications. For example, support for evolutionary design might be a better application than training, but the use of a KBS to support general design is still in the realm of research.

Given the distinction made earlier knowledge bases and knowledge based systems, it becomes clear that the very term Design Knowledge Capture is ambiguous. Design is not a domain which can be modelled; instead, we have design expertise of differing levels of generality which can be utilized in various applications. We might therefore build a knowledge based <u>system</u> to support our design activities. But we also refer to the result of these design activities as the "design" of our system. In building a DKC knowledge base we certainly will want to capture as much of this design <u>data</u> as possible in our domain model to support many different applications such as fault diagnosis, operations planning, and even evolutionary design. Additionally, much of that expertise acquired during system development which we want to retain consists of knowing how to design systems such as this one, and why this one was designed as it was. With regard to DKC, it appears that the best way to resolve these ambiguities is to consider design data, and rationale behind specific design decisions, as elements of the domain model, and thus of the knowledge base. More general design expertise should be considered as components of a knowledge based system for design, even though it was applied in this specific design instance.

With these clarifications, the DKC goal stated earlier should be more easily understood. The nature of DKC makes it a tremendous, but not insurmountable, challenge to our fledgling knowledge engineering capabilities. Much work remains to be done to develop a workable methodology and adequate tool set for DKC. NASA has set itself the goal of being a leader in the development and application of this technology. There are a number of reasons why this technology is especially critical to NASA [4], several of which will be discused below.

Motivations for Developing a Methodology for DKC

As Plato recommended in his "Republic", concepts are often best understood by examining them in the context of systems where they are applied on a large scale. Many of the systems developed by NASA represent a vast effort by hundreds of people encompassing dozens of technical disciplines. By considering what motivates NASA to pursue the development of DKC methodology, it should become clear that many of these motivations also apply in smaller projects. This discussion will therefore focus on the rationale which has led NASA to commit itself to major DKC efforts in near-term, large scale projects such as Space Station Freedom.

The importance of intelligent or knowledge based systems for NASA missions has been explicitly recognized on several occasions, both within and without the agency. Perhaps the most striking was the enactment of Public Law 98-371, which stated a requirement that:

> "The Administrator shall establish an Advanced Technology Advisory Committee in conjunction with NASA's Space Station program and that the Committee shall prepare a report by April 1, 1985, identifying specific space station systems which advance automation and robotics technologies, not in use in existing spacecraft, and that the development of such systems shall be estimated to cost no less than 10 per centum of the total Space Station costs." [5]

The seriousness with which Congress approached the use of automation and intelligent systems is reflected in the goal established for NASA's System Autonomy Technology Program:

> "The general goal to establish and maintain NASA as a world leader in this area of intelligent autonomous systems for aerospace applications will be achieved by significantly advancing the required technologies, by validating these technologies in operational environments, and by developing and maintaining world-class technical expertise, facilities and tools within the NASA organization," [6]

The Systems Autonomy Technology Program (SATP) Plan has identified development of a system knowledge base as "the central, most important technology development area" and points out that "this process must start during the design phase, where the final design represents a first baseline of factual information from which factual knowledge for the system knowledge base can be extracted". [7]

Development of intelligent systems to support NASA missions has, so far, been a matter of building individual expert systems to support particular applications. This is a very expensive process for many reasons, including the scarcity of personnel trained in knowledge engineering. The most costly activity is that of knowledge acquisition; capture of the required expertise from several sources such as documentation, data products, interviews with design and test engineers, etc. There are, as mentioned earlier, a great number of applications which could be supported by knowledge based systems utilizing this expertise. Taking one example, consider the effort required to develop a fault diagnostic expert system for each of the major components for a system if we treat each technical discipline as requiring a separate KBS. There will be enormous duplication of effort if we build a electrical fault diagnostic expert system, thermal fault diagnostic expert system, mechanical fault diagnostic expert system, etc. Furthermore, the benefits of using an integrated analysis (where thermal data are used to help diagnose an electrical problem, for example) would be totally lost. When you next consider that there are many other applications which are desirable to support our missions that could also use much of the same expertise, such as command planning, power load scheduling, evolutionary design, and so forth, then it becomes very clear that an organized, consistent approach to knowledge acquisition is both more effective and more efficient. To maximize these benefits, knowledge acquisition should be done as part of a DKC effort integrated with the development process itself. It is simply not possible to maintain the "standing armies" of engineers now used to support short term missions such as shuttle flights or Spacelab missions in an era where fifteen to thirty year operational phases are planned. The recognition of this fact has led NASA to initiate DKC efforts on its major new programs.

Development of a Methodology for DKC in NASA

In 1987 NASA initiated the Hubble Space Telescope Design/Engineering Knowledgebase (HSTDEK) Project. The primary goal of the HSTDEK Project is to enable major NASA projects to capture the design/engineering expertise they have acquired during the development of their systems in a knowledge base capable of supporting multiple applications. In order to accomplish this, current knowledge engineering technology must be extended in several areas, the new technology must be validated, and a mechanism established for transferring it to users within NASA. Six specific objectives have been identified for the project:

1. Develop a methodology for constructing multi-application, large-scale knowledge bases.
2. Develop a methodology for acquiring knowledge from multiple domain experts representing different technical disciplines, and integrating it into a single knowledge base.
3. Develop an approach for integrating knowledge engineering into the traditional engineering activitiesof a system development effort.
4. Validate this new technology in the context of a major NASA program: construction of a deep, comprehensive knowledge base for the Hubble Space Telescope.
5. Develop an in-house knowledge engineering capability for NASA to apply this new technology and support its validation.
6. Establish a program for making this new technology available to major new NASA projects, beginning with Space Station and AXAF.

This is the major effort within NASA to develop and put in place a DKC methodology. It is a joint effort between Marshall Space Flight Center (MSFC) and Ames Research Center. The basic research described in the first objective is being pursued by the Knowledge Systems Laboratory at Stanford University. Knowledge engineers at Lockheed and MSFC are the primary researchers on the second objective, while the third objective is the focus of a grant with the University of Alabama in Huntsville. The last three objectives are being handled as MSFC internal activities. HSTDEK is a five year project, funded by the Office of Aeronautics and Space Technology at NASA Headquarters. As shown in the objectives listed above, it includes research elements as well as a demonstration of DKC technology in the Hubble Space Telescope domain, which is expected to be of direct benefit to the HST during Orbital Verification of the system, and also during long term operations.

There exists a high degree of similarity in the development process for most major NASA projects. It is expected that the methodology for design knowledge capture developed in HSTDEK and validated in the Hubble Space Telescope domain will be immediately applicable in support of Space Station, AXAF, and other major NASA projects. The long operational lifetimes of these projects require DKC to meet operational objectives. In addition to retaining the design and engineering expertise developed on these projects, which usually dissipates rapidly during their operational phase, a systems knowledge base incorporating this expertise will greatly facilitate the construction of knowledge based systems for multiple applications. The problem becomes that of enhancing the knowledge base with respect to that particular application rather than starting from scratch each time. The enhanced system autonomy and productivity of ground/flight crew will result in improved mission efficiency, an extension of our capability to achieve mission goals, and an improved probability of mission success. The projected benefits to the HST Program reflect a limited subset of the expected payoffs of this project, and will now be discussed.

The planned operational lifetime of the HST is fifteen years. Even if design data can be maintained over that period by current manual/electronic means, the design and engineering expertise which provides the context for that data cannot. The HSTDEK knowledge base will provide a vehicle for making that expertise available to HST personnel throughout its lifetime. As part of its technology validation effort, HSTDEK will produce two knowledge based systems which will support specific HST applications: HSTORE and GESST. The HST Operational Readiness Expert (HSTORE) will support the Orbital Verification mission at MSFC immediately following launch of the HST with regard to telemetry monitoring and fault diagnosis of the Electrical Power System and the Pointing Control System, as well as planning for a potential Maintenance and Refurbishment (M&R) mission. The use of this system to reduce the MSFC manpower required for one and a half years of limited operational support is being investigated. This is estimated to result in a 50% decrease in the contractor support required for this purpose (approximately five man years reduction), as well as a faster response to anomalies. The Ground-based Expert System for Space Telescope (GESST) will support HST operations at GSFC, adding support for the Data Management System, Instrumentation and Control System, Mechanisms and Structures, and Thermal Control System, as well as scheduling, training and design applications. The operational impact of GESST has not yet been evaluated, but should certainly reduce GSFC reliance on MSFC and contractor experts, expedite operational support, and possibly reduce manpower required for operations. A contractor developing an AXAF proposal has unofficially estimated that the staffing per shift for that program can be reduced from 10 to 3 by use of HSTDEK technology. Considering the backlog of astronomers wishing to use the HST, the use of GESST to quickly diagnose (or even forecast and prevent) failures could result in an effective increase in that most valuable commodity, observing time, by reducing down-time. Similarly, an enhanced capability to plan maintenance and refurbishment missions should result in their occurring less frequently, which is a cost-savings in both dollars, shuttle time, and HST observing time. Another payoff from HSTDEK will be the HST knowledge base itself, which will be made available to AXAF and other NASA programs as a design aid. In addition to the HSTDEK technology deliverables per se, ten MSFC engineers will be given the training and experience needed to develop knowledge based systems. As they leave HSTDEK to work in other projects, a significant payoff will be their enhanced ability to use this technology. Coupled with the exposure gained from use of knowledge based systems to support the HST, a very high visibility NASA program, this technology transfer mechanism should result in a new generation of sophisticated knowledge based systems in NASA. HSTDEK has already generated an increased awareness of the potential of knowledge based systems in MSFC management, and more interest in applying this technology to NASA domains in the AI community.

This increased awareness has directly resulted in the incorporation of a requirement for DKC in the Phase C/D Request for Proposal issued for AXAF. Both bidders on that contract addressed the DKC issue in their proposal, and the winner (TRW) will now proceed to implement their DKC program. It is inappropriate to discuss specifics at this point since negotiations a re still in progress, but the presence of even a limited DKC element in AXAF planning is a very positive step for the technology, and a validation of its perceived importance in NASA.

By far the greatest challenge, and the strongest requirement, for DKC in NASA is the development of Space Station Freedom. In response to the Congressional direction discussed above, NASA prepared a document titled Process Requirements for Design Knowledge Capture, whose objective is "to define design knowledge capture requirements placed on the work package contractors in support of the Space Station design community." [8] This document goes on to say that DKC is intended to include both design objects and designer's knowledge. The response of bidders for the work package contracts in the area of DKC varied greatly in both the depth and scope of their proposed approaches. The Space Station Program Office has given the responsibility for consolidating procedures and data relating to DKC to the Systems Engineering and Integration Information Planning Group. This group includes members from the Program Office, the Work Package Centers, and the International Partners. One current objective of the group is to revise the Process Requirements Document for DKC, and provide a coordinated input to the Space Station Program Requirements Document. Unfortunately, the planning for DKC at that level is now focused on use of the Technical Management Information System to capture text oriented design information, rather than integration of knowledge engineering tools with the traditional engineering activities. Discussions with the Space Station Strategic Plans and Programs Division as part of a recent Advanced Automation Study [9] has resulted in the definition of a three year task beginning in Fiscal Year 1989 to assess the ability of the Software Support Environment (SSE) workstations which will be used for development of all Space Station operational software to support the DKC activity. This task will

> "Build upon on-going Design Knowledge Capture (DKC) and use SSE Workstations as mechanism for performing DKC activities; define DKC capabilities (and information content requirements) by working with a fully designed spacecraft and assess the potential of the SSE Workstation and its software to perform portions of SS DKC; identify hardware/software modifications required to the SSE Workstation to support DKC."[10]

This task will apply and extend the experience gained in HSTDEK in order to develop a plan for initiating a full DKC effort for Space Station. This approach will treat knowledge based systems as elements of the planned SS operational software set.

NASA is therefore aggressively pursuing the development of DKC methodology and planning for its incorporation in major new programs, including the Hubble Space Telescope, Advanced X-Ray Astronomical Facility, and Space Station Freedom. Given the magnitude of the design and engineering efforts required to develop these systems, and their projected fifteen to thirty year operational lifetimes, there is simply no practical way to complete their missions without the use of Design Knowledge Capture.

References

1. Barr, A., and Feigenbaum, E. A. (Eds.) The Handbook of Artificial Intelligence, Volume One, William Kaufman, Inc., Los Gatos CA, 1981, Page 3.

2. Freeman, M. S., "An Investigation of Multi-User Expert Systems", Dissertation, University of Alabama in Huntsville, May 1987.

3. "Systems Autonomy Technology Program (SATP) Plan" (FY89 Draft), NASA/Ames Research Center, August 1988, page A-160.

4. Freeman, M. S., and Hooper, J. W., "Factors Affecting the Development of Expert Systems in NASA", First Conference on Artificial Intelligence for Space Applications, NASA/Marshall Space Flight Center and the University of Alabama in Huntsville, October 1985.

5. "National Aeronautics and Space Administration Research and Development, 98 Stat. 1227, Research and Program Management Report" Public Law 98-371, July 1984.

6. "Systems Autonomy Technology Program (SATP) Plan" (FY89 Draft), NASA/Ames Research Center, August 1988, Foreword.

7. "Systems Autonomy Technology Program (SATP) Plan", NASA TM 100999, NASA/Ames Research Center, December 1987, Page 16.

8. "Process Requirements for Design Knowledge Capture", JSC 30471, Johnson Space Center, December 15, 1986, Page 1-2.

9. "Space Station Advanced Automation Study Final report", Strategic Plans and Programs Division, Office of Space Station, NASA Headquarters, May 1988.

10. Program Operating Plan (POP) 88-2 Guidelines for the Space Station Transition Definition Program, Office of Space Station, NASA Headquarters, May 16, 1988, Page 2-11.