Task-Based Knowledge Management*

D. Leake
Computer Science
Lindley Hall
Indiana University
150 S. Woodlawn Ave.
Bloomington, IN 47405
leake@cs.indiana.edu

L. Birnbaum & K. Hammond
Intelligent Information Laboratory
Computer Science Department
Northwestern University
1890 Maple Avenue
Evanston, IL 60201
{birnbaum,hammond}@ils.nwu.edu

C. Marlow
Computer Science
University of Chicago
1100 East 58th Street
Chicago, IL 60637
cameron@ils.nwu.edu

H. Yang
PDC, MD 270, GB-D68
Ford Motor Company
21500 Oakwood Blvd.
Dearborn, MI 48124
hyang1@ford.com

Abstract

Case-based reasoning is receiving much attention as a technology for building knowledge repositories that can be queried for task-relevant information. Taking the CBR problem-solving model seriously, however, suggests the value of a much stronger integration between knowledge management systems and the tasks that they serve. In this integrated view, knowledge management systems should be designed to do *just-in-time retrieval,* anticipating task-based information needs and satisfying them automatically before the user requests information, and should learn unobtrusively by monitoring the user's task performance. Key issues include how to integrate knowledge access into the user's problem-solving process, how to automatically provide the user with task-relevant information from multiple sources, and how to build up knowledge for transmission between task phases and for long-term storage. This paper describes how these issues are addressed in the Stamping Advisor, a system to aid the design of stamped automotive parts. This system automatically presents the designer with needed information in a natural way, uses CBR and task-focused information retrieval to access useful information, and automatically captures relevant information to support downstream task processes and build its memory of cases.

Introduction

The goal of knowledge management is to manage repositories of information to facilitate information access and reuse. This may involve providing tools that a user can apply to access information resources, such as data warehouses, or developing practices that encourage employees to contribute to knowledge repositories (O'Leary 1998). CBR techniques provide a basis for representation and retrieval of experiential information (Becerra-Fernandez & Aha 1999), and have already been successfully applied for large-scale corporate knowledge sharing (Kitano & Shimazu 1996; Klahr 1997), as well as for local knowledge capture and reuse in many contexts.

We believe that the integrated problem-solving process suggested by the CBR problem-solving model can have additional impact for the design of knowledge management systems. In current knowledge management systems, knowledge capture and access are task-motivated but normally require the user to make the link between the information and the task. The user must decide that he or she needs to seek information, must determine the type of information to seek to further the task, and must select tools to access knowledge repositories to seek that information. After the task is complete, the user must take extra steps to document the task process for capture. In contrast, the CBR problem-solving model treats knowledge access and capture as automatic parts of the problem-solving process (Kolodner 1993). This view suggests a model of knowledge management in which knowledge access and capture are fully integrated into support for the user's task. Central questions for this model are how to track ongoing knowledge use, how to automatically provide the information needed at each point, and how to incrementally build up and transmit knowledge to support further processing.

This paper illustrates the task-based approach with the *Stamping Advisor*, a knowledge management system for design support. The Stamping Advisor was developed in consultation with domain experts at the Ford Motor Company, which is now applying its methods to new systems. The Stamping Advisor is designed around a model of the design task and the process by which that task is carried out. The system exploits that model to proactively present the right information at the right time and to use cases not only for knowledge storage, but also for transmission of knowledge to downstream tasks. This approach provides two benefits. The first benefit is *just-in-time* information access: the system anticipates the user's needs instead of requiring the user to specify them, generates queries

*This research is supported in part by the Ford Motor Company under award 0970-355-A200. David Leake is currently a Visiting Professor at the Intelligent Information Laboratory at Northwestern University, on sabbatical leave from Indiana University, and he thanks the Laboratory and the Northwestern Computer Science Department for their support. His work is also supported in part by NASA under award NCC 2-1035.

to relevant information sources, and provides the information in a format that meshes with the user's own conceptualizations of the task at hand. The second benefit is that, as the system tracks information needs, it can incrementally collect useful information to build up case representations both for long-term storage and as a medium for communicating information to downstream tasks. This facilitates the current task processes and supports the growth of system knowledge.

Supporting Feasibility Assessment

Task-based knowledge management requires an analysis of the tasks to support and the information products they produce. Our testbed domain is the design of stamped metal parts (e.g., hoods and fenders) for automobiles. The design of such parts depends on a combination of styling considerations and practical considerations (such as aerodynamics and structural requirements) for the finished part, and practical constraints on stamping particular shapes.

In the part design process, manufacturing engineers interact with design engineers to ensure the manufacturing feasibility of the parts being designed. Manufacturing engineers do a feasibility study to assess potential difficulties in the formability of the parts and other design issues, and evaluate projected costs to improve production efficiency and ensure the best possible design. Their feasibility judgments are made mainly on the basis of part geometry and material properties. For example, some shapes are difficult to make because of the stretchability of metal; some flange shapes will result in unpredictable "springback" of the metal towards its original shape after stamping. A feasibility engineer identifies problems and proposes part modifications to the design engineer, who revises the design or justifies the problematic decisions in an iterative cycle.

Feasibility engineers report that prior experiences often play an important role in their judgments. Consequently, a primary focus of our system is to provide the feasibility engineer with access to prior experiences, both to help in locating problems and to help in finding solutions. At the start of the project, the Ford Motor Company had already captured paper records of issues and decisions and had gathered initial "seed cases" in a database.

The Stamping Advisor supports the analysis of a design by first retrieving similar previous parts that suggest potential issues, and providing the part descriptions, issue descriptions, and descriptions of their resolutions to the feasibility engineer. It uses the information about prior issues to automatically generate queries to retrieve web-based information relevant to those issues, searching on-line Ford manuals. Its final product is the report that the feasibility engineer provides downstream at the end of feasibility analysis.

Fitting Case Presentation to Reasoning

In order to integrate information presentation with user problem-solving, task-based knowledge management systems must provide their information in a way that fits the user's reasoning process. Interviews with feasibility engineers established that one of their reasoning styles is to sequentially scan the CAD image of a stamped part, moving around the border looking for potential issues. As they scan the image, they are often reminded of problems affecting similar parts in the past. The Stamping Advisor augments their memories: Given a new part to analyze, it automatically retrieves cases for parts with similar characteristics. The indexing process is described in (Leake *et al.* 1999).

The system presents case information by superimposing issue summaries on the CAD image being examined, grouping cases according to the regions they concern. To provide a rapid indication of problem areas, summaries are color-coded according to the severity of issues they involve, with green for regions where prior cases support feasibility, yellow for limited problems, and red for more serious problems. Figure 1 shows an issue summary screen.

To examine the cases relevant to a particular region, the engineer can click on the summary box to see the associated case information. As shown in Figure 2, a window appears with information about issues in prior parts that have similar regions. These can be selected to view how they were resolved.

Focusing Search for Extra Information

The Stamping Advisor also exploits the task context to automatically generate queries for useful information to supplement the case library and help explain the significance of case features. For example, when feasibility engineers examine problems and determine fixes, they refer to Ford's engineering rules and on-line "best practices" documents (e.g., for style guidelines and for ways to reduce manufacturing costs). One of the goals of the Stamping Advisor project is to provide a framework for satisfying task-based information needs from traditional on-line sources as well as cases. It does this extracting relevant features of the current design and prior problem cases and using them to generate task-relevant queries to a standard search engine indexing on-line references. We have indexed Ford's on-line manuals with the commercial search engine Verity, but the approach applies to any search engine.

When the engineer selects a potential problem from a retrieved case to examine, the Stamping Advisor extracts information from the current design to describe the part being designed, and from the problem case to describe the issue being considered. This information includes the type of car, type of part, and region currently of interest, as well as the problem pointed out by the retrieved case. It uses this information to automatically generate a focused query aimed at retrieving information relevant to the type of part and problem.

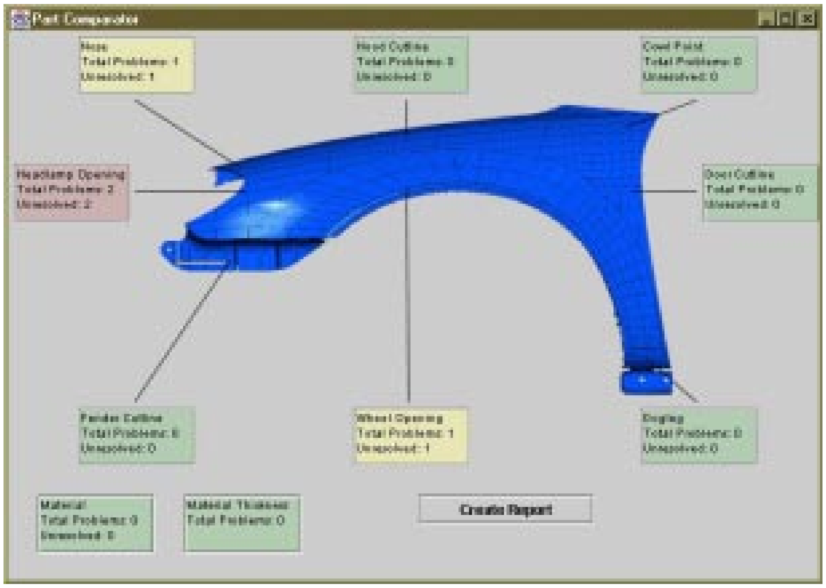

Figure 1: Screen image from the Stamping Advisor's issue summary screen for a fender design.

The system presents the query to the user, as shown at the bottom of Figure 2. The user can initiate search, based on the system-generated description of the part involved, of the problem involved, or of both. The user may also edit the query as desired (e.g., to compare the problem with the same type of problem affecting a different line of car). The query generation capability enables the system to automatically supplement case knowledge with task-relevant information from additional networked knowledge sources.

Integration Across Tasks

In addition to providing information to support the feasibility engineer's task, the Stamping Advisor supports the incremental capture and sharing of information across tasks. In our view, each design phase should automatically access information about the previous steps in the design process to provide context for current reasoning, and should produce products that can be used by the reasoning processes downstream. The downstream process from feasibility assessment is evaluation of the final design. At the close of the feasibility assessment process, the system generates a Final Report Document to aid other design team members who need to understand or evaluate the feasibility engineer's work. In our model of the evaluator's task, downstream team members need to know (1) the part being examined, (2) the issues considered, (3) how they were disposed of, and (4) the prior parts providing evidence relevant to the issues and decisions. Prior to this system, all this information was collected manually. The information in parts 1, 2, and 4 is now collected automatically as the Stamping Advisor is used. The information in 3 is entered by the user, requiring no more effort than the effort already involved in the manual process of writing out suggestions for changes.

Case Capture and Transmission

The information in the Final Report Document is used for two purposes. The first is knowledge transmission: it is the information product of the feasibility analysis process and is sent on automatically to the evaluation phase. The second is to add to the case library: It contains the information needed for a new case.

In general, we believe that knowledge management systems should build cases by establishing a growing record, in which information is added at each step and the cumulative record is made available to each downstream process. For the part design process, for example, information used to characterize part designs in the CAD system (e.g., model, year, and part number, and a pointer to the CAD file), provides an initial record of the design. The next step, feasibility analysis, identifies issues and revisions to add to this record. Finally, the evaluation step adds overall results. This could be augmented with the capture of relevant information from the shop floor as well (e.g., reporting unanticipated problems). This provides a natural means for adding to the case library. In addition, if most of case capture is automatic and user inputs are controlled (e.g., though menus), the representation of cases can be standardized, increasing the reliability of future retrievals.

Principles from the Project

The Stamping Advisor illustrates five principles for integrating CBR into design support tools, and, more

Figure 2: Window providing access to prior cases, issues, and resolutions for a headlamp opening problem.

generally, for using CBR in task-based knowledge management (Leake *et al.* 1999):

- Seamless interaction: Interaction with the combined system must parallel the user's own problem-solving process, to make the information available without leaving the problem-solving context.
- Proactive support: The system must automatically provide the right information when it is needed, rather than placing the burden on the user to formulate requests.
- Integration with other knowledge sources: The system must provide a task-based facility to gather relevant information across all available information resources, supplementing case information with other types of information as appropriate to the task.
- Integration across tasks: The system must support not only the immediate reasoning task but also the downstream tasks it serves. The system should automatically access information about the previous tasks to provide a context for its reasoning.
- Experience capture: Each processing episode must generate new cases or update old ones to capture information for future reasoning and downstream processing.

Perspective

Case-based reasoning has already received considerable attention in design support systems (see Kolodner (93) for an overview), including systems that provide references to prior problems in order to warn users about potential problems to avoid (e.g., (Domeshek & Kolodner 1992)). The Stamping Advisor, however, aims for a tighter integration between the CBR component and the overall task supported by design. It also places more control of retrieval in the hands of the user, in a closer spirit to conversational CBR (Aha & Breslow 1997).

CBR is also receiving increasing attention as a technology for knowledge repositories and knowledge reuse (Becerra-Fernandez & Aha 1999; Klahr 1997), with particular recent focus on facilitating access to existing textual repositories of information (Lenz & Ashley 1998). A difficult problem is how to provide focused retrieval, a traditional strength of CBR, when using unstructured information sources. Rissland and Daniels (1995) address this by using traditional retrieval of structured cases to provide relevant examples that seed the relevance feedback mechanism of an IR system that searches a larger pool of textual cases. Our system integrates CBR and IR by extracting important problem features from the task context and prior cases, and using those features to automatically generate task-relevant queries for a standard search engine.

The importance of analyzing information use and information flow in organizations is well-known in the knowledge management community. However, to our knowledge, previous systems do not provide the integrated proactive information access illustrated by the Stamping Advisor.

Conclusion

This paper presents a task-based approach to knowledge management that provides just-in-time retrieval, anticipating information needs and providing task-relevant information within the task context, before it is requested. The paper illustrates this approach with the Stamping Advisor system, a support system for engineering design. The Stamping Advisor reflects a model of the design process and exploits that model to anticipate information needs and present useful information in a natural way. The system also uses its task model to supplement case information with additional information gathered by task-focussed IR techniques and to determine useful information to transmit downstream. This forms a basis for automatically augmenting system knowledge by unobtrusively collecting important information during the task process.

References

Aha, D., and Breslow, L. 1997. Refining conversational case libraries. In *Proceedings of the Second International Conference on Case-Based Reasoning*, 267–278. Berlin: Springer Verlag.

Becerra-Fernandez, I., and Aha, D. 1999. Case-based problem solving for knowledge management systems. In *Proceedings of the Twelfth Annual Florida Artificial Intelligence Research Symposium.* Menlo Park: AAAI. In press.

Domeshek, E., and Kolodner, J. 1992. Towards a case-based aid for conceptual design. *International Journal of Expert Systems* 4(2):201–220.

Kitano, H., and Shimazu, H. 1996. The experience sharing architecture: A case study in corporate-wide case-based software quality control. In Leake, D., ed., *Case-Based Reasoning: Experiences, Lessons, and Future Directions.* Menlo Park, CA: AAAI Press. 235–268.

Klahr, P. 1997. Knowledge management on a global scale. In Gaines, B.; Musen, M.; and Uthurusamy, R., eds., *Proceedings of the 1997 Spring Symposium on Artificial Intelligence in Knowledge Management*, 82–85. Stanford, CA: AAAI.

Kolodner, J. 1993. *Case-Based Reasoning.* San Mateo, CA: Morgan Kaufmann.

Leake, D.; Birnbaum, L.; Hammond, K.; Marlow, C.; and Yang, H. 1999. Integrating information resources: A case study of engineering design support. In *Proceedings of the Third International Conference on Case-Based Reasoning.* Berlin: Springer Verlag. In press.

Lenz, M., and Ashley, K., eds. 1998. *Proceedings of the AAAI-98 workshop on textual case-based reasoning.* Menlo Park, CA: AAAI Press.

O'Leary, D. 1998. Enterprise knowledge management. *Computer* 54–61.

Rissland, E., and Daniels, J. 1995. Using CBR to drive IR. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 400–407. Montreal: IJCAI.