Integrating Information Resources: A Case Study of Engineering Design Support*

David B. Leake,[1] Larry Birnbaum,[2] Kristian Hammond,[2] Cameron Marlow,[3] and Hao Yang[4]

[1] Computer Science Department, Lindley Hall, Indiana University,
150 S. Woodlawn Ave, Bloomington, IN 47405, U.S.A., leake@cs.indiana.edu
[2] Intelligent Information Laboratory, Computer Science Department,
Northwestern University, 1890 Maple Avenue, Evanston, IL 60201, U.S.A.,
{birnbaum,hammond}@ils.nwu.edu
[3] Computer Science Department, The University of Chicago,
1100 East 58th Street, Chicago, IL 60637, U.S.A., cameron@ils.nwu.edu
[4] PDC, MD 270, GB-D68, Vehicle Operations, Ford Motor Company, 21500
Oakwood Blvd., Dearborn, MI 48124, U.S.A., hyang1@ford.com

Abstract. The development of successful case-based design aids depends both on the CBR processes themselves and on crucial questions of integrating the CBR system into the larger task context: how to make the CBR component provide information at the right time and in the right form, how to access relevant information from additional information sources to supplement the case library, how to capture information for use downstream and how to unobtrusively acquire new cases. This paper presents a set of design principles and techniques that integrate methods from CBR and information retrieval to address these questions. The paper illustrates their application through a case study of the Stamping Advisor, a tool to support feasibility analysis for stamped metal automotive parts.

1 Introduction

An experienced designer's memory of prior design experiences can be a powerful aid during the design process. When the designer who faces a new task is reminded of similar previous tasks, those remindings may suggest related solutions and warn of potential problems to avoid. Case-based design support systems leverage this process: They augment the designer's own memory by providing relevant cases from a library of prior experiences.

* This research is supported by the Ford Motor Company under award No 0970-355-A200. David Leake is currently a Visiting Professor at Northwestern University, on sabbatical leave from Indiana University, and thanks the Intelligent Information Laboratory and the Northwestern Computer Science Department for their support. His research is supported in part by NASA under award No NCC 2-1035. Copyright ©1999 Springer–Verlag. This paper appears in the *Proceedings* of ICCBR-99 but is not camera-identical to the proceedings version. This paper is available from the archive at `http://www.cs.indiana.edu/~leake/INDEX.html`.

Case-based design has long been an active area of case-based reasoning research, and numerous case-based design aids have been implemented to support a wide range of design tasks (see (Kolodner, 1993) for some examples of these systems). Fully realizing the benefits of such systems, however, requires addressing additional issues beyond the case-based design support process itself. In order to maximize the usefulness of case-based design aids, they must be designed not as stand-alone systems but as integral parts of a single unified framework that supports all phases of the design process and the multiple actors that are often involved, and that draws on the multiple available information resources. Developers of such systems must address crucial questions of integrating the CBR system into the larger task context: how to make the CBR component provide information at the right time and in the right form, how to exploit other information sources in concert with case information, and how to capture information for use downstream and to unobtrusively acquire new cases. This paper presents a set of design principles and techniques addressing these questions. It illustrates their application through a case study of the Stamping Advisor, a tool to support feasibility analysis for stamped automotive parts.

2 The Stamping Advisor Domain

Automotive body design is a crucial task in automobile development. Body design has a profound impact on the vehicle's appeal and function, and the body is the most expensive component of the vehicle to manufacture. Stamped body parts, which make up the major portion of the body subsystem, are designed under constraints arising from aesthetic considerations, structural and functional requirements, cost concerns, and the availability of manufacturing resources.

Body styles are developed in an iterative process between the designers and feasibility engineers who examine the design for potential manufacturing issues. These include formability issues, which may result in splitting or wrinkling of the metal after the stamping process; manufacturing process complexity issues, such as shapes that must be stamped with a large number of dies (increasing costs), and quality concerns due to material properties and feature shapes, which may add significant cost to die testing or affect the quality or consistency of the final product.

The feasibility engineer's task is to identify potential problems, to justify why they are likely to occur, to estimate the costs that will be incurred if they are not addressed, and to propose design revisions to remedy them. Feasibility engineers report that they often base their judgments on specific experiences with prior designs. However, new engineers begin their work without this library of experiences, and even experienced engineers may not have had experience with the most relevant designs for a particular problem. Multiple information resources exist to aid the feasibility analysis task, such as records of experiences with prior designs, stored in paper and electronic forms. However, it may be difficult or excessively time-consuming for engineers to locate the needed information. Likewise, communicating their decisions and justifications is often cumbersome:

The standard method for communicating their decisions downstream is to fill out and send a paper form.

Key questions for improving this process are how to provide better access to experiences and other engineering knowledge, and how to improve the usefulness of the information when it is reapplied. A collaboration was established between the Intelligent Information Laboratory at Northwestern University and the Vehicle Operations and Visteon divisions at the Ford Motor Company to investigate integrated case-based design support systems to address these questions. The company already had captured paper records of feasibility assessment issues and decisions, some of which had been placed in a database, providing a library of seed cases. The research question was how, given a set of feasibility analysis cases and the standard manuals used by feasibility engineers, to access and present them to maximize their usefulness to the design process.

Thus one goal of the project was *information integration* (Knoblock and Levy, 1998): to develop methods for satisfying the designer's information needs using cases and other information sources, for integrating the CBR system to automatically produce the information needed downstream, and for supporting unobtrusive case acquisition from available information. The Intelligent Information Laboratory developed the *Stamping Advisor* system, described in this paper, to demonstrate a framework for this design support process, and its approaches are now being applied to new systems at the Ford Motor Company.

3 Principles for Integrated Intelligent Design Support

The Stamping Advisor system embodies five general principles for the integration of case-based design support systems into the design environment. These principles are:

- **Seamless interaction**: Interaction with the combined system must parallel the feasibility engineer's own problem-solving process.
- **Just-in-time retrieval**: The system must proactively anticipate information needs and 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 link all available information resources, presenting prior cases, supplementary information to help understand the cases or apply their lessons, and additional information as appropriate to the task.
- **Integration across tasks**: The system must serve 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, and should produce products that can be used by the reasoning processes downstream.
- **Experience capture**: Each processing episode must provide new cases in a usable form.

These principles are related to basic tenets of the case-based reasoning cognitive model (Kolodner, 1994; Leake, 1998; Schank, 1982): That accessing and storing cases is a natural part of task performance and that models of knowledge access must reflect the task context. Our design support framework extends these principles to anticipate the user's needs, accessing relevant information wherever it is available, and extends the target of support beyond the current user to capture and transmit relevant information downstream.

3.1 Realizing these principles

Achieving a design support system that respects the previous principles requires addressing a number of CBR issues. Integrating the system with the feasibility engineer's reasoning and providing just-in-time support requires modeling his or her reasoning process, and especially modeling when and why particular cases and other information resources are retrieved. Integrating multiple knowledge sources depends both on appropriate task-based indexing and on methods for similarity assessment and retrieval that can be applied to preexisting documents and other information sources that differ from traditional cases. Experience capture depends on methods for case acquisition. The remainder of this paper discusses how the Stamping Advisor system addresses each of these issues.

4 Coordinating Case Presentation with the Reasoning of Feasibility Engineers

One of the goals of the Stamping Advisor project was to make case presentation fit the engineer's own reasoning. This is done in two ways: by designing the case presentation interface to fit the engineer's reasoning style, and by using knowledge of the engineer's task to anticipate the engineer's information needs and provide information proactively.

Feasibility engineers are given a computer-generated image of the part to evaluate, produced by the computer-aided design (CAD) system that the engineer used to generate the design. Interviews with feasibility engineers established that one of their reasoning styles is to sequentially scan the image, tracing around the boarder of the part looking for portions of the design that raise feasibility issues. The primary system interface provides a CAD image of a part, with different regions annotated by information about relevant cases. This makes it easy for the engineer to follow his or her normal process of scanning the design.

Given a design whose feasibility needs to be determined, the system presents a summary of the cases retrieved and the issues involved, using a graphical display of a part image with annotations concerning the number of issues found for each region of the part and their resolutions. The graphical interface organizes case information geometrically according to the regions of the part. For each region, it provides a summary of the cases found that involve issues for that region. The summaries of the issues for each area of the design are highlighted with color-coded warnings to identify the most problematic regions (green when surrogates support feasibility, yellow for limited problems, red for more serious problems). Figure 1 shows the issue summary interface for an automobile fender. In the screen display, the leftmost box, describing the headlamp opening, is highlighted in red because previous cases identified two potential issues that could not be resolved. The boxes for the nose (upper left) and wheel opening (bottom center) are highlighted in yellow, because each one includes one unresolved problem. No other problems were found, so all other boxes are highlighted in green.

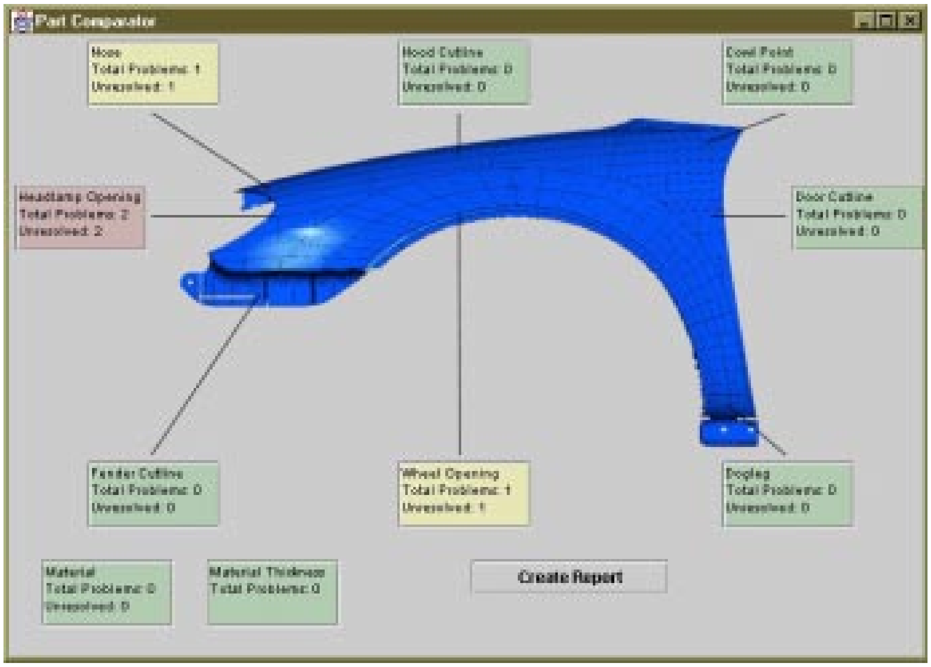

Nose
Total Problems: 1
Unresolved: 1

Hood Cutline
Total Problems: 0
Unresolved: 0

Cowl Point
Total Problems: 0
Unresolved: 0

Headlamp Opening
Total Problems: 2
Unresolved: 2

Door Cutline
Total Problems: 0
Unresolved: 0

Fender Cutline
Total Problems: 0
Unresolved: 0

Wheel Opening
Total Problems: 1
Unresolved: 1

Dogleg
Total Problems: 0
Unresolved: 0

Material
Total Problems: 0
Unresolved: 0

Material Thickness
Total Problems: 0

Fig. 1. Screen image from the Stamping Advisor's issue summary screen.

To see additional information for a region, the feasibility engineer clicks on the boxes for the displayed issue sets to select a region of interest. A window appears with information about problems in prior parts (called "surrogate parts") in which that region was similar. The engineer can select problems from this list to see how they were resolved. In some cases, the design will have been revised to repair the problem, suggesting a possible revision to consider. In others, the previous engineer may have detected mitigating factors that were originally overlooked, which made the problem inapplicable; these suggest factors for the engineer to check in the current design. In some cases, the prior engineer may have decided that the problematic design feature was so valuable aesthetically that it counterbalanced the extra production costs; in that situation the old case contains information about the estimated costs to consider when weighing whether to allow the potential problem to remain. The interface for this process is shown in Figure 2.

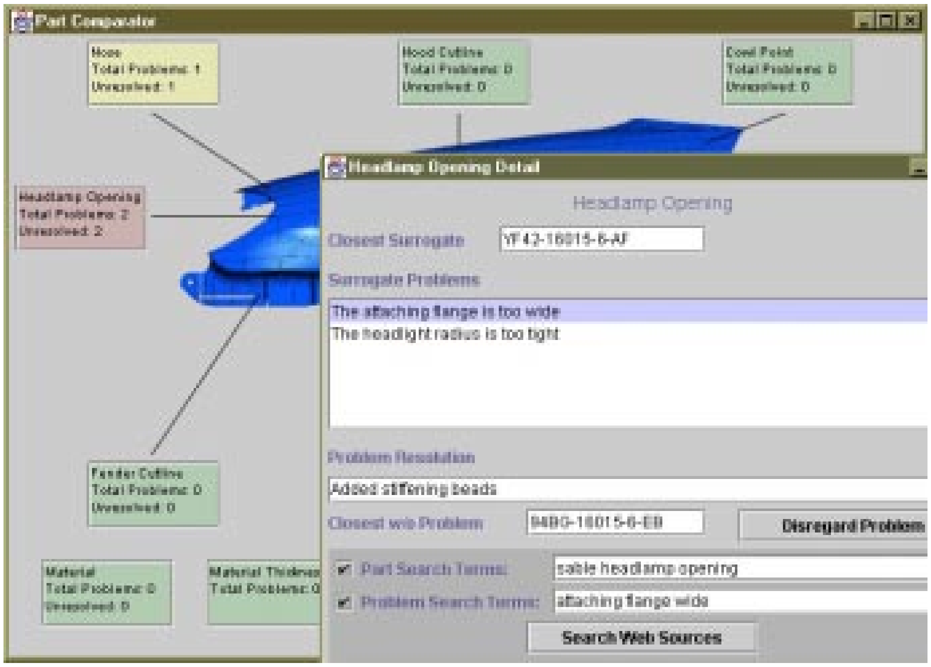

Hood
Total Problems: 1
Unresolved: 1

Hood Cutline
Total Problems: 0
Unresolved: 0

Cowl Point
Total Problems: 0
Unresolved: 0

Headlamp Opening
Total Problems: 2
Unresolved: 2

Fender Cutline
Total Problems: 0
Unresolved: 0

Material
Total Problems: 0
Unresolved: 0

Material Thickness
Total Problems: 0

Closest Surrogate: YF42-16015-6-AF

Surrogate Problems:
The attaching flange is too wide
The headlight radius is too tight

Problem Resolution: Added stiffening beads

Closest w/o Problem: 94BG-16015-6-EB

☑ Part Search Terms: sable headlamp opening
☑ Problem Search Terms: attaching flange wide

Fig. 2. Presentation of relevant surrogates, issues, and resolutions.

5 The Case Retrieval Process

In the Stamping Advisor, each part type is associated with a predefined set of classes of features to examine for feasibility. For a fender, there are ten such classes. Eight of these are associated with geometric regions of the part (e.g., the class of features involved in the headlamp opening), while two concern characteristics of the material used (e.g., stamping aluminum parts instead of steel parts involves special feasibility issues concerning sheet metal thickness).

When the system retrieves cases for potential issues, candidate cases are filtered according to the type of part being analyzed; for example, when examining the feasibility of a fender, only prior experiences with fenders are considered for retrieval. Within the cases for the given type of part, the system retrieves one set of relevant cases for each class of features to examine. For example, the system retrieves cases for fenders with similar headlamp openings to suggest feasibility issues associated with the design of the headlamp opening; it retrieves cases for fenders with similar wheel opening tabs to suggest feasibility issues associated with the design of the wheel opening tabs, and so forth. After cases have been filtered by the part and the type of part feature under consideration, the basic matching process is a nearest-neighbor algorithm using feature weightings developed for the domain.

In some instances, relationships may exist between distinct classes of features, so that simply considering the regions independently is not sufficient. For example, one stamping problem is "springback," in which a panel returns to its original shape after stamping (e.g., because of the amount of stretching required and the material used). The amount of "springback" may depend on the relationships between the shapes of two adjacent regions. In such cases, the relationship across types of features is recorded and used to adjust the weighting of retrieved cases. For example, if both adjacent regions have features that suggest springback, the weight of the cases suggesting springback is increased compared to the weights that were derived from looking at each region alone before considering the supporting relationship between them.

6 Integrated Information Access

Cases are helpful for warning of potential problems and suggesting prior solutions. However, additional information may be needed to assess the relevance of prior issues, to determine the applicability of old solutions, or to develop new solutions reflecting changed constraints. For example, Ford maintains on-line manuals with design recommendations for keeping stamping costs reasonable and for maintaining consistent styling. Given that these information sources will often be required to supplement retrieved cases, access to this information is important.

Keeping with the philosophy of integrating the CBR system, our goal was to use knowledge of the user's task and task context to automatically guide the search for this information: to automatically present the engineer with the supplementary information that is useful, given the knowledge that it is being retrieved in response to specific issues in a specific case. To provide this support, the Stamping advisor uses tracking information about the current task to automatically formulate targeted queries that can go against documents indexed by standard search engines. The delivered system demonstrates this capability by automatically generating queries to retrieve relevant style guidelines from the Ford Advanced Feasibility Guidelines for Styling.

6.1 Query Generation and Document Retrieval

As a product of the manual feasibility analysis process, textual information such as part names, part numbers, problem descriptions, feature names, and the vehicle name are recorded in a paper description. This information has been encoded into the database from which the cases are retrieved, and consequently is available for every part handled by the Stamping Advisor. This text is sufficient to distinguish parts at a textual level.

The Stamping Advisor uses this descriptive information, combined with its model of task relevance, to form queries to other information resources. Specifically, when a feasibility engineer is considering a feature, the system automatically forms queries to gather additional information about related features or problems from on-line resources. Four pieces of information establish the context for this query: the names of the vehicle, part, and problematic feature, and the textual description of the problem in question. These are extracted form the record of the current design. The system removes words contained in a standard stop list and makes a query from the remaining terms.

For example, when the feasibility engineer examines the headlamp opening problems highlighted in Figure 1, one of the issues is that the attaching flange is too wide. The Stamping Advisor generates a query containing “Sable headlamp opening” for the part under consideration, and “attaching flange wide” for the problem. Upon the feasibility engineer’s request, this query is used to search for relevant guidelines in on-line manuals. Before initiating search, the engineer can request that the query be focused on only similar parts or similar problems, and can edit the query text as desired (e.g., to replace “Sable” to compare the styling on a different line of car). The query presentation interface is shown at the bottom right of Figure 2.

Once created, this query can be passed to any typical Internet search engine to search selected resources. In our implementation, we use the document indexing system Verity to index documents such as the on-line Ford Style Guide illustrated in Figure 3. Verity processes queries by stemming each of the given words, broadening the search to other possible forms of the terms, and assigning a numerical score. This score is based first on the number of word matches and then on the density of those matches within a given document. The list of matches is presented to the feasibility engineers, who can select documents to retrieve.

7 Integration Across Tasks

Previous case-based design support tools have a natural goal: aiding a designer in his or her task. However, in industrial settings, the designer’s task is only one step in an extended process. For example, in stamping design, one or more designers initially formulate the design, a feasibility engineer critiques the design and makes suggestions, and the design is refined though an iterative cycle of changes and critiques. When a design is finalized, downstream design team members may need to evaluate the design, its potential issues and the designers’ justifications for why they matter (or do not matter), and how they were resolved. Ideally, design aids should support this entire *process* rather than supporting only one individual step. This requires the sharing of information across tasks.

A tenet of our design support principles is that the design support system for any particular task should automatically access information about the previous tasks to provide a context for its reasoning, and should produce products that can be used by the reasoning processes downstream. Work is under way on augmenting the CAD system used for initial design to automatically capture the specification information used in feasibility analysis cases (e.g., to capture the part number, part type, vehicle, and a pointer to the CAD file), to be passed

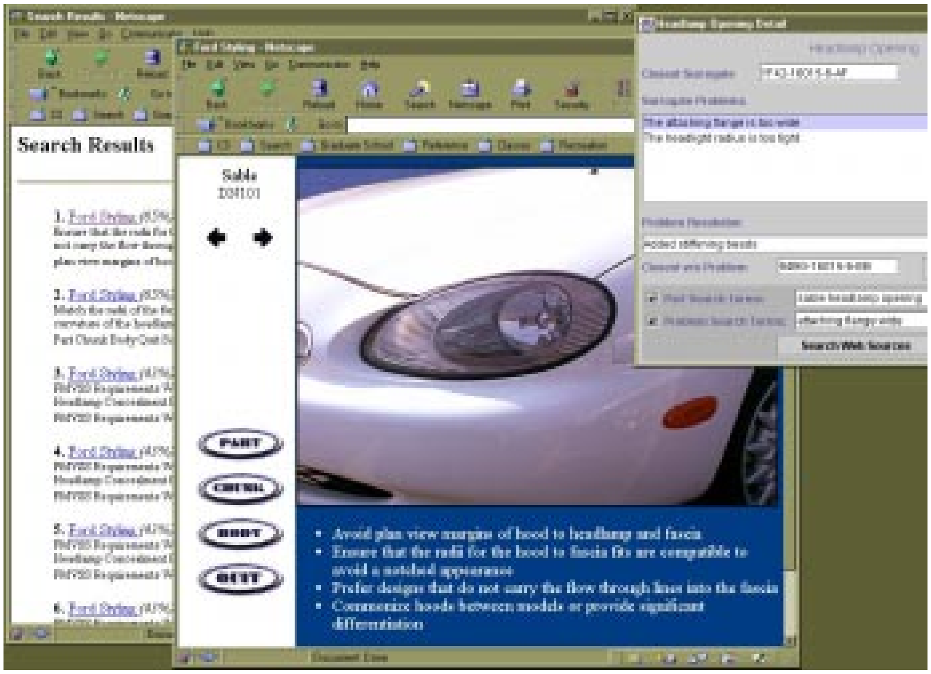

- Avoid plan view margins of hood to headlamp and fascia
- Ensure that the radii for the hood to fascia fits are compatible to avoid a notched appearance
- Prefer designs that do not carry the flow through lines into the fascia
- Commonize hoods between models or provide significant differentiation

Closest Surrogate: 1F43-16015-B-AF

Surrogate Problems:
The attaching flange is too wide
The headlight radius is too tight

Problem Resolution: Added stiffening beads

Closest w/o Problem: 6G93-16016-B-BB

☑ Part Search Terms: sable headlamp opening

☑ Problem Search Terms: attaching flange wide

Fig. 3. Style guide page retrieved as relevant to the problem of headlamp opening being too tight.

automatically to the Stamping Advisor at the start of feasibility analysis. This will provide additional integration between the task of the initial designer and the feasibility engineer.

At the close of the feasibility assessment process, the system generates a Final Report Document to aid upstream or downstream design team members who need to understand or evaluate the feasibility engineer's work, replacing documentation generated by hand. In our model of the evaluation task, the information needed is: (1) the part being examined, (2) the issues considered, (3) how they were disposed of, and (4) the surrogates providing evidence relevant to the issues and decisions. A sample Final Report Document is shown in Figure 4.

8 Case Capture

Ford maintains an extensive library of reports of feasibility analysis problems and solutions in paper form. However, as is often the case in applying CBR, there is a bottleneck in translating this information into a usable case form. The ability of the Stamping Advisor to create Final Report Documents suggests a way to alleviate this bottleneck. In the Stamping Advisor, a user's decisions about appropriate surrogates, the problems they predict, and the ultimate disposition

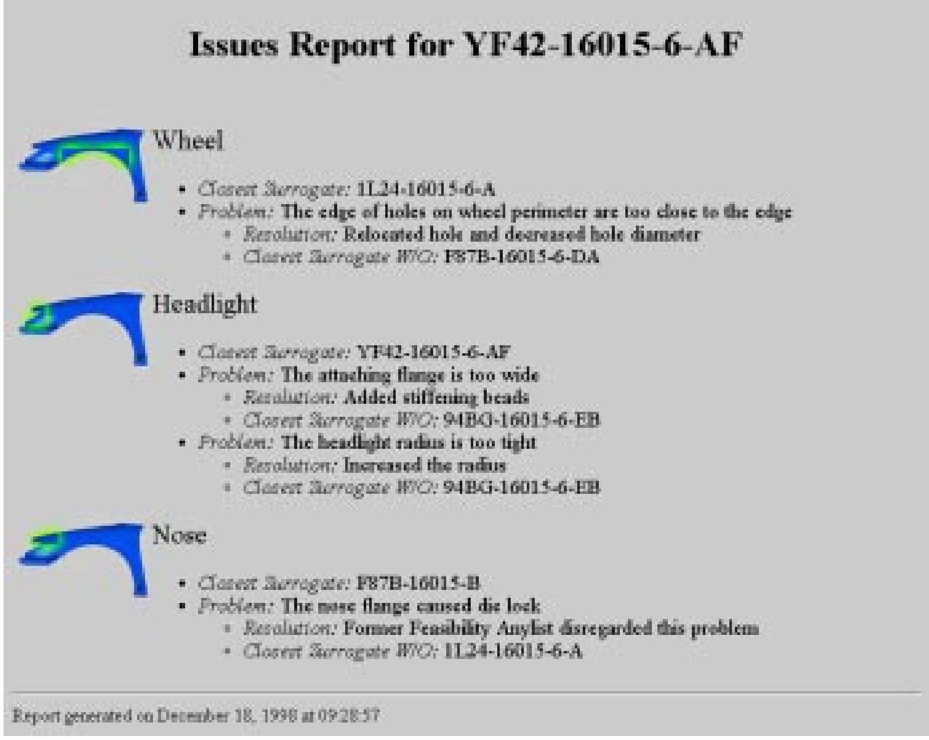

Issues Report for YF42-16015-6-AF

Wheel

- *Closest Surrogate:* 1L24-16015-6-A
- *Problem:* The edge of holes on wheel perimeter are too close to the edge
  - *Resolution:* Relocated hole and decreased hole diameter
  - *Closest Surrogate W/O:* F87B-16015-6-DA

Headlight

- *Closest Surrogate:* YF42-16015-6-AF
- *Problem:* The attaching flange is too wide
  - *Resolution:* Added stiffening beads
  - *Closest Surrogate W/O:* 94BG-16015-6-EB
- *Problem:* The headlight radius is too tight
  - *Resolution:* Increased the radius
  - *Closest Surrogate W/O:* 94BG-16015-6-EB

Nose

- *Closest Surrogate:* F87B-16015-B
- *Problem:* The nose flange caused die lock
  - *Resolution:* Former Feasibility Anylist disregarded this problem
  - *Closest Surrogate W/O:* 1L24-16015-6-A

Report generated on December 18, 1998 at 09:28:57

Fig. 4. Final report document.

of the problems are captured by the system during the feasibility assessment process. These are used to create the Final Report Document. This document is produced as the by-product of the user's decision-making and does not require additional effort on his or her part beyond that already required to convey the needed information downstream. This document automatically combines information captured from the user with other background information, gathering all the information needed to generate a new feasibility assessment case.

This case capture framework gathers data when they are available at each phase of the design process, not just during feasibility analysis. The growing record is made available to each downstream process for reasoning from existing data and addition to the record. In particular, information is built up during initial part design, feasibility analysis, and final decision-making on how to proceed on a part.

Information used to characterize part designs in the CAD system (e.g., (model, year, and part number, and a pointer to the CAD file) provide an initial record of the design. Current seed cases include geometric features, and work is ongoing to support the addition of geometric features to new cases. Ideally, general-purpose automatic geometric matching procedures (e.g., (Coulon and

Steffens, 1994)) could be applied to the designs. However, given the specialized domain and comparatively small number of important features, special-purpose feature extraction routines also appear practical. Some of these have been developed by Ford. Alternatively, because the engineer must already document the important geometric features when describing problems to generate the downstream report, it would be comparatively simple to tag these features according to a predefined vocabulary of standard features that can then be used for matching.

When the Final Report Document is provided electronically to the person who determines the final disposition of the request, that person can enter the final decision to complete the case information. By controlling the information that can be entered at each step of the process (e.g., though menus), cases can be standardized. However, the ability to do textual searches provides the additional capability to search through free-form comments, etc.

In summary, our framework integrates case capture across different parts of the design process and uses cases as a vehicle both for sharing knowledge as it is gathered and for long-term knowledge capture. In particular, case content should:

1. Be built up incrementally as a natural part of the problem solving process.
2. Be used incrementally during the process, as soon as it has been generated.
3. Provide a full record of relevant information at the end of the process, in the form needed for future use by tools to support feasibility assessment.

This supports rapid growth of case information and the standardization of provided information.

9 Relationship to Previous Work

9.1 Case-Based Design Support

A wide range of case-based design support tools has been developed for numerous tasks such as architectural design (Gebhardt *et al.*, 1997; Goel *et al.*, 1991; Hua and Faltings, 1993; Maher *et al.*, 1995; Smith *et al.*, 1995), conceptual design of aircraft subsystems (Domeshek *et al.*, 1994; Leake and Wilson, 1999), autoclave layout design (Hinkle and Toomey, 1995), device design (Goel, 1989; Sycara *et al.*, 1991), and circuit design (Vollrath, 1998). The Stamping Advisor's task is most closely related to that of the load validator in the system Clavier (Hinkle and Toomey, 1995), which warns users about potential problems in new autoclave layouts by presenting users with similar prior layouts and their outcomes. A crucial issue in autoclave layout design is the interacting effects of components of the layouts, and these interactions are hard to explain and separate. Consequently, Clavier based its predictions on the similarity of the previous layouts, taken as a whole, with entire current designs. In the Stamping Advisor domain, problems can be localized by the feasibility engineer. Consequently, Stamping Advisor cases represent problems at the level of the individual regions they affect (with additional checks for interactions that span multiple regions), which facilitates transfer of problem information to new contexts (for example, headlamp opening problems can be predicted based on prior experiences with the headlamp openings in very different styles of fenders). The Stamping Advisor also differs in using cases not only to advise, but also to capture and communicate the rationale underlying design decisions taken in response to its advice.

The Stamping Advisor demonstrates a number of principles for integrating CBR into the engineering design process. First, the system brings CBR into the feasibility engineer's normal reasoning process by integrating case-based support with the CAD tools already used to create and examine designs for stamped parts. This approach is similar to those taken by the FABEL (Gebhardt *et al.*, 1997) and CADRE (Hua *et al.*, 1996) projects, both of which integrate the CBR system with existing CAD systems. It differs, however, in using a very specific task model to automatically determine the types of information to provide and when to provide it with just-in-time retrieval. In contrast, FABEL provides a "virtual construction site" that the engineer can navigate, and a tool kit from which the designer selects tools to perform particular types of retrievals. The Stamping Advisor uses its model of how the feasibility analysis task is done to anticipate specific information needs and proactively determine what information is needed and how to retrieve it.

9.2 Integrating CBR and IR

The Stamping Advisor also goes beyond case-based support to integrate multiple knowledge sources. There is considerable current interest in the use of CBR for textual cases, and in the use of information retrieval methods to access them (Lenz and Ashley, 1998). A challenging question is how to maintain the strengths of CBR—the pragmatic focus that traditional CBR provides—while exploiting the generality of IR methods for assessing the similarity of documents. This depends on bridging the gap between task-relevant indexing used in CBR and methods that can be applied to unstructured textual data. (Rissland and Daniels, 1996) present one method for this integration in the retrieval of legal cases. Their system first performs a feature analysis to do a traditional CBR retrieval of the most relevant cases from a case library represented in a carefully structured form. It then uses the textual descriptions of those cases as seed examples for the relevance feedback mechanism of a text-based information retrieval system, which generates queries to retrieve similar texts from a larger library of textual case descriptions. The Stamping Advisor uses task-based characterizations more directly: it directly generates a search engine query from relevant problem features. Because the role of each component in the query is readily apparent, the Stamping Advisor also provides the user with the capability to revise this query before search to reflect additional information goals that may not be known to the system.

9.3 Case Capture

A crucial issue for scaling up CBR applications is knowledge capture. The Stamping Advisor system is designed to facilitate this through knowledge capture during use. Feasibility analysis is a "natural" CBR domain (Mark *et al.*, 1996), in that the manual feasibility analysis process includes extensive paper documentation for each design case. However, the primary case acquisition mode we envision is from system use itself. Even if *no* cases were available in the system case library, the system would be useful as a convenient interface for recording feasibility information (now recorded on paper) and aiding search through online resources. Thus feasibility engineers have the incentive to use the system, and their use provides cases that will increase its usefulness as sufficient data is gathered to take full advantage of the CBR component.

10 Conclusions

The Stamping Advisor project illustrates a set of principles for integrating case-based reasoning systems into the larger task context. The system was designed to provide an open architecture for case and other information retrieval based on features of the current design, and to exploit and support the flow of information from successive steps of the design process. To make the system natural to use, the interaction is designed to parallel the feasibility engineer's own problem-solving process and to automatically provide just-in-time access to the right cases, rather than placing the burden on the user to formulate requests. The system uses its task model to generate focused IR queries to access additional knowledge sources, retaining the capability for the user to adjust those queries to explore additional topics. The system does automatic knowledge capture, gathering information about each interaction and using it for a dual purpose: to provide the information needed downstream of the reasoning task and package new cases for future use.

The central lesson of this work is that the development of successful case-based design aids must depend not only on the CBR processes themselves but on crucial questions of integrating CBR system into the larger task context: making the system automatically provide information when it is needed and in the right form, accessing relevant information from additional information sources, and communicating and capturing information. We are continuing to strengthen this integration as the current system is refined. One goal, for example, is to fully integrate the Stamping Advisor into the initial CAD design process, to immediately warn the original designer of potential problems while the design is being generated. We believe that CBR fits naturally into a new mode of knowledge management that not only tracks where documents *are*, but tracks *how they are used* and *where they are needed* to access multiple information sources to provide the right information at the right time.

References

[Coulon and Steffens, 1994] C.-H. Coulon and R. Steffens. Comparing fragments by their images. FABEL Report 13, Gesellschaft für Mathematik und Datenverarbeitung mbH, 1994. Pages 36–44.

[Domeshek *et al.*, 1994] E. Domeshek, M. Herndon, A. Bennett, and J. Kolodner. A case-based design aid for conceptual design of aircraft subsystems. In *Proceedings of the Tenth IEEE Conference on Artificial Intelligence for Applications*, pages 63–69, Washington, 1994. IEEE Computer Society Press.

[Gebhardt *et al.*, 1997] F. Gebhardt, A. Voß, W. Gräther, and B. Schmidt-Belz. *Reasoning with complex cases.* Kluwer, Boston, 1997.

[Goel *et al.*, 1991] A. Goel, J. Kolodner, M. Pearce, and R. Billington. Towards a case-based tool for aiding conceptual design problem solving. In R. Bareiss, editor, *Proceedings of the DARPA Case-Based Reasoning Workshop*, pages 109–120, San Mateo, 1991. DARPA, Morgan Kaufmann.

[Goel, 1989] A. Goel. *Integration of Case-Based Reasoning and Model-Based Reasoning for Adaptive Design Problem Solving.* PhD thesis, The Ohio State University, 1989.

[Hinkle and Toomey, 1995] D. Hinkle and C. Toomey. Applying case-based reasoning to manufacturing. *AI Magazine*, 16(1):65–73, Spring 1995.

[Hua and Faltings, 1993] K. Hua and B. Faltings. Exploring case-based design - CADRE. *Artificial Intelligence in Engineering Design, Analysis and Manufacturing*, 7(2):135–144, 1993.

[Hua *et al.*, 1996] K. Hua, B. Faltings, and I. Smith. CADRE: Case-based geometric design. *Artificial Intelligence in Engineering*, 10:171–183, 1996.

[Knoblock and Levy, 1998] K. Knoblock and A. Levy, editors. *Proceedings of the AAAI-98 workshop on AI and information integration.* AAAI Press, Menlo Park, CA, 1998.

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

[Kolodner, 1994] J. Kolodner. From natural language understanding to case-based reasoning and beyond: A perspective on the cognitive model that ties it all together. In R. Schank and E. Langer, editors, *Beliefs, Reasoning, and Decision Making: Psycho-Logic in Honor of Bob Abelson*, pages 55–110. Lawrence Erlbaum, Hillsdale, NJ, 1994.

[Leake and Wilson, 1999] D. Leake and D. Wilson. Combining CBR with interactive knowledge acquisition, manipulation and reuse. In *Proceedings of the Third International Conference on Case-Based Reasoning*, Berlin, 1999. Springer Verlag. In press.

[Leake, 1998] D. Leake. Cognition as case-based reasoning. In W. Bechtel and G. Graham, editors, *A Companion to Cognitive Science*, pages 465–476. Blackwell, Oxford, 1998.

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

[Maher *et al.*, 1995] M. Maher, B. Balachandran, and D. Zhang. *Case-based reasoning in design.* Erlbaum, Hillsdale, NJ, 1995.

[Mark *et al.*, 1996] W. Mark, E. Simoudis, and D. Hinkle. Case-based reasoning: Expectations and results. In D. Leake, editor, *Case-Based Reasoning: Experiences, Lessons, and Future Directions*, pages 269–294. AAAI Press, Menlo Park, CA, 1996.

[Rissland and Daniels, 1996] E. Rissland and J. Daniels. The synergistic application of CBR to IR. *Artificial Intelligence Review*, 10:441–475, 1996.

[Schank, 1982] R.C. Schank. *Dynamic Memory: A Theory of Learning in Computers and People*. Cambridge University Press, Cambridge, England, 1982.

[Smith *et al.*, 1995] I. Smith, C. Lottaz, and B. Faltings. Spatial composition using cases: IDIOM. In *Proceedings of First International Conference on Case-Based Reasoning*, pages 88–97, Berlin, October 1995. Springer Verlag.

[Sycara *et al.*, 1991] K. Sycara, R. Guttal, J. Koning, S. Narasimhan, and D. Navinchandra. CADET: a case-based synthesis tool for engineering design. *International Journal of Expert Systems*, 4(2):157–188, 1991.

[Vollrath, 1998] I. Vollrath. Reuse of complex electronic designs: Requirements analysis for a CBR application. In P. Cunningham, B. Smyth, and M. Keane, editors, *Proceedings of the Fourth European Workshop on Case-Based Reasoning*, pages 136–147, Berlin, 1998. Springer Verlag.