

Towards More Creative Case-Based Design Systems

Linda M. Wills and Janet L. Kolodner

College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
linda@cc.gatech.edu, jlk@cc.gatech.edu

Abstract

Case-based reasoning (CBR) has a great deal to offer in supporting creative design, particularly processes that rely heavily on previous design experience, such as framing the problem and evaluating design alternatives. However, most existing CBR systems are not living up to their potential. They tend to adapt and reuse old solutions in routine ways, producing robust but uninspired results. Little research effort has been directed towards the kinds of situation assessment, evaluation, and assimilation processes that facilitate the exploration of ideas and the elaboration and redefinition of problems that are crucial to creative design. Also, their typically rigid control structures do not facilitate the kinds of strategic control and opportunism inherent in creative reasoning. In this paper, we describe the types of behavior we would like case-based design systems to support, based on a study of designers working on a mechanical engineering problem. We show how the standard CBR framework should be extended and we describe an architecture we are developing to experiment with these ideas.¹

Introduction

Creativity in design derives from enumerating several solution alternatives, redescribing and elaborating problem specifications, and evaluating proposed solutions, based on criteria and constraints that go beyond the stated constraints on a solution. It arises out of a confluence of processes (including problem redescription, remembering, assimilation, and evaluation), which interact with each other in complex ways. Often creativity arises from interesting strategic control of these processes, which in themselves may be quite mundane (Boden 1990, Chandrasekaran 1990, Gero & Maher 1993, Navinchandra 1992).

These processes rely heavily on previous design experiences and knowledge of designed artifacts (Goel & Chandrasekaran 1992, Hinrichs 1992, Kolodner & Penberthy 1990, Kolodner & Wills 1993). An expert de-

signer knows of many design experiences, accumulated from personally designing artifacts, being given case studies of designs in school, and observing artifacts designed by others. The designer draws on these experiences to perform such activities as generating design alternatives, reformulating and elaborating the problem specification or proposed solutions, and predicting the outcome of making certain design decisions. The experiences that are most valuable are often those that are highly contextualized pieces of knowledge about artifacts, such as how a device behaves in some context of use, circumstances in which it can fail, and knowledge about situations that might come up not only in use, but in all phases of its life cycle.

Given the nature of these experiences, we believe case-based representations and reasoning techniques lend themselves to supporting creative design. Research in case-based reasoning (CBR) has provided extensive knowledge of how to reuse solutions to old problems in new situations, how to build and search case libraries (for exploration of design alternatives), and how to merge and adapt cases. It has developed powerful techniques for partial matching and the formation of analogical maps between seemingly disparate situations (Kolodner 1993).

However, most existing CBR systems are not living up to their potential. They tend to adapt and reuse old solutions in routine ways, producing robust but uninspired results. They do not attempt to extend their exploration by deriving constraints and preferences that improve or go beyond those stated in the original problem. (See (Kolodner 1993, appendix) for a recent survey.)

Some of this potential is buried in processes that have been downplayed or even missing in most standard CBR systems. In particular, little research effort has been directed towards the kinds of situation assessment, evaluation, and assimilation processes that facilitate the exploration of ideas and the elaboration and redefinition of problems that are crucial to creative design. Also, to facilitate the kinds of opportunism inherent in creative reasoning, CBR systems need to break out of their typically rigid control struc-

¹This research was funded in part by NSF Grant No. IRI-8921256 and ONR Grant No. N00014-92-J-1234.

ture to allow flexible interleaving and communication among processes. In addition, more research attention must be paid to the strategic control mechanisms that guide a creative designer in deciding what to do next.

In this paper, we describe the types of behavior we would like case-based design systems to support, based on an exploratory study of designers working on a mechanical engineering problem. We show how the standard CBR framework should be extended and we describe an architecture we are developing to experiment with these ideas. We end with a set of open issues.

What Do Creative Designers Do?

To gain insights into the knowledge and reasoning involved in creative design, we observed a four-person team engaged in a seven-week undergraduate mechanical engineering (ME) design project. The task was to design and build a device to quickly and safely transport several eggs from one location to another. The device could be constructed from any material, but its size, weight, and cost were restricted.

After exploring several schemes for launching, moving, stopping, and protecting the eggs, the team decided to use a cylindrical egg carrier (of radius 7 cm., length 22.5 cm.), with the eggs wrapped in pipe insulation to protect them inside the carrier. The carrier was dropped down (0.8 m.) from a starting platform and would roll into a target zone (within a 5 m. radius of the starting platform). The team had two possible launch mechanisms up until the final design demonstration day: a spring mechanism and a simple ramp (the spring launch base could be inverted to become a ramp, which was the final choice). In both cases, a string, with one end attached to the launch base, was wrapped around the device, so that as the cylinder dropped, it spun down the string, hit the ground, and rolled into the target zone. The wrapped string gave the carrier momentum and it also prevented it from rolling beyond the target zone.

One of us participated as a member of the team, allowing us to become immersed in the issues and to observe the design process in a natural setting, in both informal and "official" team meetings. We recorded the group's conversations on audiotapes and collected copies of all their design documents and drawings.

We are particularly intrigued by a set of three processes we observed underlying many creative design activities: 1) generation of multiple descriptions or views of a problem, 2) gradual emergence of evaluative issues, constraints and preferences, and 3) serendipitous recognition of solutions to pending problems, sometimes seeing new functions and purposes for common design pieces in the process. We are not claiming that this is a complete set. (For example, our design study has revealed a variety of influences on creativity from collaborative activity.) Rather, we are interested in these processes because they are key processes in design that current case-based systems neglect.

Problem Redescription. The initial problem statement given to our designers was ambiguous, incomplete, contradictory, and underconstrained. They spent a great deal of effort to turn it into something with more detail, more concrete specifications, and more clearly defined and consistent constraints. An important part of this process involved attempting to understand the problem, view it from multiple perspectives, and redescribe it in terms familiar to the designers. They had to refine and operationalize several vague or abstract constraints, while sometimes having to abstract constraints that were too specific.

For example, many of the ideas of one designer, who had a keen interest in automobiles, came from recalling devices and concepts from the car domain, such as shock absorbers, unit-body vs. single-frame construction, and air-bags. Being able to recall these required viewing the problem of protecting the eggs as one of absorbing shock or transferring energy and as a problem of protecting passengers in general, not just eggs.

Our designers also explored the given constraints, deliberately stretching or strengthening them to see what ideas became possible. For example, the initial problem statement was ambiguous about whether or not the device could land (i.e., touch down) short of the target zone and then move into it. The designers considered the extreme possibility of landing as far short of this zone as possible, in which case the device would not fly at all, but would be pushed off or lowered to the ground, where it would then move itself into the safety zone. Visualizing this possibility reminded them of devices, such as elevators and yo-yo's, that could implement parts of this behavior.

This continual elaboration and redescription of the problem helped the designers derive connections between the current problem and similar problems in other domains, facilitating cross-contextual transfer of design ideas. It also primed them to serendipitously recognize relevant objects in the environment that might be reused for a new purpose.

Evaluation. One of the key forces driving evolution of the problem specification is the evaluation of proposed design alternatives. Evaluative issues emerge in the course of evaluating. Designers do not merely depend on constraints that have already been specified. Rather, they bring up additional constraints and criteria as proposals are examined. Proposed solutions often remind them of issues to consider. The problem and solution "co-evolve" (Fischer 1993).

One interesting criteria that emerged in the course of the ME design project was *versatility* – the ability of the device to apply in more than one situation. This criteria was not mentioned or required in the original statement of the problem. It arose in response to ambiguity in the initial problem statement, which described three similar problems but did not specify which one would be assigned. Each problem differed only in the device's starting position (from either the center of a

child's wading pool or from a platform of one of two heights) and in its target destination distance. (This is similar in the real world to situations in which the engineers are designing for multiple potential customers with different needs). To deal with the uncertainty and reduce the complexity this variability introduced, the designers began searching for solutions that could be used to solve all three problems or could be easily adapted to apply to each. That is, they began to evaluate proposals on the basis of versatility in addition to the other criteria already in the problem specification. Being able to do this is central to creative design.

Assimilation. Problem redescription provides not only a means for recalling relevant solution alternatives, but also a vocabulary for describing and, in many cases, reinterpreting objects in the designer's environment. This often leads to a new way of viewing the function of some object and facilitates the recognition of potential solutions to pending problems in the external environment.

For example, our designers went to a home improvement store for materials for a spring launch mechanism. While comparing the strengths of several springs by compressing them, they noticed that the springs tended to bend. One designer wrapped a hand around the spring to hold it straight as it was compressed and said the springs would each need to be enclosed in a tube to keep them from bending. Another added that the tube would need to be collapsible (to compress with the spring). The designers could not think of an existing collapsible tube and did not want to build one due to time pressure. They gave up on the springs and started thinking about egg protection. During their search for protection material, they walked through the bathroom section of the store, where they saw a display of toilet paper holders. They immediately recognized them as collapsible tubes which could be used to support the springs.

By playing with the springs, noticing problems and suggesting fixes, the designers formed a specific, concrete, and operationalized description of what a solution would look like to the bending-springs problem. However, the toilet-paper holder was not recalled on the basis of this description. Instead, the description was used to reinterpret the toilet paper holder when it was encountered in the external environment and to recognize its additional function of preventing springs from bending upon compression. The designers were able to interpret objects seen in the environment, or recalled from memory, from a new viewpoint. This viewpoint was based on descriptions and feature dimensions that had been revealed to be important in attempts to solve recent and pending problems.

We refer to this process as *assimilating* the objects into a problem context. It not only involves reinterpreting solution alternatives under consideration, but also comparing and contrasting alternatives with one another, along the dimensions relevant to the problem

context. This helps reveal those that are not really new ideas, so that they can be ignored. It can also cause new evaluative issues to emerge as new dimensions or criteria are generated to distinguish seemingly identical ideas.

Strategic Control. The designers we observed did not follow a rigid, methodical plan detailing what to do next. Rather, they moved fluidly between various problem pieces and design processes (e.g., idea generation, adaptation, critiquing, problem refinement, elaboration, and redefinition) in a flexible and highly opportunistic manner.

Our designers employed a variety of strategic control heuristics, some of which are opportunistic. For example, when an alternative was proposed that satisfied some desired criteria extremely well compared to the other alternatives, they directed their efforts toward elaborating that alternative, optimistically suspending criticism or discounting the importance of criteria or constraints that were not satisfied as well. Sometimes this led to reformulation of the problem as constraints were relaxed or placed at a lower priority.

Being able to take advantage of such opportunities requires being able to judge whether progress was being made along a certain line of attack and to choose which ideas are more promising or more likely to lead to something unusual and novel.

Some strategic control heuristics are more deliberate, based on reflection. For example, one heuristic our designers used was to try quick, easy adaptations of a proposed solution first before stepping back and reformulating the problem or relaxing constraints. Other deliberate heuristics attempted to make non-standard substitutions, apply adaptation strategies in circumstances other than the ones they were meant for, and merge pieces of separate solutions with each other in nonobvious ways.

In many cases, the processes that are composed together leading to a novel idea are not in themselves novel and may be quite mundane. The trick is knowing when to do them.

How CBR Systems Can Do Better

Most current CBR systems tend to stick to well-known interpretations of problems and routine ways of adapting old solutions, neglecting exploration of alternatives if something good enough has been found. We believe the CBR paradigm can be extended to support more creative problem solving.

Problem Redescription. Problem redescription corresponds closely to the process of situation assessment – redescription of a problem in the vocabulary of the indexing system. In most CBR systems, situation assessment is skipped; the assumption is made that the initial representation of the problem is sufficient for solving the problem. But, as our observations show, investigating a problem in depth makes available a large

set of relevant cues for retrieval. Generating multiple ways of describing a problem provides several different contexts for specifying what would be relevant, if remembered.

Research on indexing has found that it is the combination of setting up a context for retrieval and having already interpreted something in memory in a similar way that allows retrieval. When some case or piece of knowledge is entered into memory, it is not always possible to anticipate how it might be used. Situation assessment processes aim to bridge that gap by helping to redescribe a new problem in a way that is similar to something seen before.

Research into situation assessment and problem reformulation (e.g., in CASEY (Koton 1988), CYRUS (Kolodner 1983), MINSTREL (Turner 1994), BRAINSTORMER (Jones 1992), and STRATA (Lowry 1987)), show different ways it can be done. However, these techniques have not yet made it into widespread use in practical CBR systems. They should certainly be included in any system aimed at reuse of experience across domains.

Evaluation. CBR systems currently evaluate solutions by checking a set of constraints that have been given to the system. Evaluative procedures are typically buried within case manipulation to predict or test whether a modified case satisfies the specified constraints. Observations of our designers suggests that evaluation should play a more prominent role in case-based design systems, allowing evaluative issues to emerge in the course of evaluating. Navinchandra (1991) calls this *criteria emergence* and shows an example of how it can arise from case-based projection. In addition to criteria, constraints in general (Prabhakar & Goel 1992) and relative priorities among them also gradually emerge. This type of evaluation is a key driving force within creative design, feeding back to situation assessment and guiding case manipulation.

Assimilation. A key idea underlying dynamic memory (Schank 1982), one of the principle foundations of case-based reasoning, is that remembering, understanding, and learning are all inextricably intertwined. The ability to determine where something fits in with what we already know (understanding) is a key part of being able to assimilate objects in our environment into our problem solving. This environment includes not only external objects, but also cases that have been retrieved, elaborated and adapted. Understanding how these fit into a problem context may involve a useful reinterpretation of something already in memory, suggesting in a new way of indexing it.

Strategic Control. Our exploratory study suggests that a linear, sequential composition of CBR processes is much too simple. In reality, these processes are highly intertwined and interact in interesting ways. For example, problem elaboration and redescription tac-

tics specify contexts for search that retrieval processes use, while evaluation of recalled or adapted alternatives feeds information back to these situation assessment tactics, resulting in even better contexts for search. In some cases, what suggests a particular problem refinement or redescription results from trying to confirm the legality of a proposed solution during evaluation and finding a loophole or ambiguity in the current problem specification. In addition, comparing and contrasting a proposed solution with other proposals during assimilation can bring new evaluative issues into focus.

CBR systems need to break out of their typically rigid control structure and allow more interaction and opportunism among processes. This requires making strategic control mechanisms explicit, so they can be easily modified, reasoned about, extended, and learned. More research needs to be directed at identifying and capturing the types of strategic control heuristics designers use.

Proposed Architecture

We are developing an experimental case-based system that emphasizes the processes of situation assessment, evaluation, and assimilation, integrating them with the usual CBR processes of retrieval, elaboration (case manipulation, adaptation, merging, prediction), and learning. It has a flexible, opportunistic control structure which allows us to keep control tactics separate, explicit, and modifiable.

The processes within our system are not applied in a strictly linear succession. Rather, the system has a blackboard-style architecture. The processes are centered around and act upon data structures that represent the evolving problem specification and the set of design alternatives under consideration.

Situation assessment procedures act on the problem specification to evolve it along multiple directions. Evaluation examines design alternatives, checking them against the current specification, to reveal inconsistencies, ambiguities, and incompleteness in the specification that suggest new redescrptions. Evaluation also brings up new criteria, and constraints which are incorporated into the problem specification.

Elaboration procedures transform alternatives under consideration into new alternatives by applying a variety of adaptation and merging strategies. These strategies are typically suggested by the critique formed by an evaluation of some alternative. Elaboration procedures also augment alternatives with information derived about their consequences and expected behavior. These "data collection" elaborations are currently accomplished by manual augmentations of alternatives with experimental data, but in general can be achieved by case-based projection, simulation, actual experimentation, or visualization.

The evolving problem description is also used by both the retrieval and the assimilation processes. Retrieval interfaces with a library of cases which models,

in part, long-term memory. The problem description is used as a probe into memory to pull relevant design cases into consideration (for evaluation, elaboration, etc.). The assimilation process is the dual of retrieval. It accumulates design alternatives proposed (i.e., those retrieved, elaborated, or viewed directly in the external environment) into the pool of design alternatives under consideration, organizing the alternatives with respect to each other.

The data structure holding the set of design alternatives forms an extension of the long-term memory. We call this extension the “problem context.” The evolving problem description determines the focal vocabulary of the current problem context. As the specification evolves, the focus changes on the relevant vocabulary to be used for organizing alternatives in the memory (e.g., shape, construction cost, personal safety). In a sense, the problem context is providing a point of view with respect to which objects in the environment and cases recalled can be interpreted and organized by the assimilation process.

The coordination of the various processes is controlled by explicit strategic control mechanisms. There are a set of monitoring procedures, associated with each of the processes, which watch for opportunities for some task to be performed. The opportunities noticed are placed on an “opportunity agenda.” Opportunities are chosen and pulled from the agenda by strategic control heuristics. For example, a monitor associated with the assimilation process watches for an alternative to be added that is much better than any other alternative proposed so far, with respect to some desired criterion. This yields an opportunity to change the problem description by increasing the priority of that criterion and/or by relaxing constraints that are not met by that proposal. This simulates the behavior of changing the relative importance among criteria to accommodate an unexpectedly good solution that is stumbled upon. An example strategic control heuristic would be to pursue elaboration opportunities for alternatives that satisfy a desired criteria extremely well before pursuing evaluative processes that would negatively critique the alternatives. This simulates the behavior of optimistically pursuing an idea, suspending all but constructive criticism.

Status, Limitations and Open Issues

Our system currently has implemented procedures for evaluation, assimilation, and retrieval, as well as data structures representing the case library, pool of design alternatives, evolving problem specification, and the opportunity agenda data structure. We have standard agenda management routines. However, these routines currently do not model the ephemeral nature of opportunities (which can either expire or be forgotten). Several monitors surrounding the assimilation process have been implemented, but we still need to define and capture those relevant to the other processes.

Much more work is needed to identify and define strategic control heuristics, situation assessment procedures, and elaboration techniques. Also, not all strategic control mechanisms are triggered by noticing an opportunity. Some may become applicable due to some complex condition that must be inferred through reflection. (For example, realizing that you are reasoning in circles might cause you to make an effort to try a brand new technique.) More research needs to focus on how to represent and infer these kinds of conditions and also how the application of these more reflective strategic control mechanisms can be interleaved with the triggering of opportunistic ones.

We are starting to understand how criteria, constraints, preferences, etc., emerge during evaluation, but more effort is needed in modeling this emergence.

There are a number of interesting open issues concerning how assimilation is managed when the design problem is complex, having several interacting sub-problems, each of which have different sets of alternatives and requirements. Assimilation must find the appropriate problem context for interpreting and evaluating a given design alternative. The ability to do this facilitates the serendipitous recognition of solutions to pending problems, as we saw in the bending-springs problem. (See also (Seifert et al. 1994).)

Another open issue is that the designers we studied were not expert mechanical engineers. An interesting empirical question is: would experts, having knowledge of “design principles,” behave differently? It may not be the expert vs. novice distinction, but how open-ended the problem is, that is important. After all, the students were familiar with and experienced in solving everyday mechanical problems using objects in their world. We believe that for open-ended, nonroutine problems, expert designers are likely to display the same sorts of behaviors as do our students.

Finally, there are some aspects of creative design that we have not yet explored. In particular, we would like to analyze more carefully the influences collaboration had on creativity in the design project. Our agenda-based model of opportunity management lends itself to simulating the exploration of several opportunities in parallel, and employing multiple control strategies at once. This will allow us to simulate these aspects of collaborative activity and use computational experiments to explore hypotheses about the role of collaboration in creative design.

Conclusion

Our intention in building our system is not to automate design, but to test our hypotheses about the cognition of creative design. We are trying to understand creative processes better, using a case-based cognitive model. As we increase our understanding (and in the process, push CBR technology), we will be able to answer the question how best to assist human designers. This may include 1) aiding the formalization, reformu-

lation, and refinement of specifications (Reubenstein & Waters 1991, Johnson, Benner, & Harris 1993), 2) bringing up evaluative issues (Domeshek & Kolodner 1993), 3) retrieving pending problem contexts to help recognize the applicability of solutions, or 4) proposing new control strategies.

We are taking a case-based approach to understanding creative design for two reasons. One is that many creative design activities are highly memory-intensive and rely on past design experiences, so case-based reasoning has much to offer in this study. The other is that we hope to make case-based systems themselves more creative. By using the paradigmatic tools CBR provides, we are starting to find computational models of the behaviors and processes we observed in our exploratory study. At the same time, our modeling attempts have deepened our understanding of case-based processes and memory issues and have suggested extensions that will yield more creative design systems in the future.

Acknowledgements

We appreciate the insightful discussions we have had with Terry Chandler, Eric Domeshek, Lucy Gibson, Todd Griffith, Kenneth Moorman, Nancy Nersessian, Ashwin Ram, and Mimi Recker. We would like to thank Otto Baskin, Jon Howard, and Malisa Sarntinoranont, for their invaluable cooperation. We also appreciate the insights and helpful comments of our anonymous reviewers.

References

Boden, M. 1990. *The Creative Mind: Myths and Mechanisms*. New York, NY: Basic Books.

Chandrasekaran, B. 1990. Design Problem Solving: A Task Analysis. *AI Magazine*. 11(4): 59-71.

Domeshek, E.A., and Kolodner, J.L. 1993. Using the points of large cases, *Artificial Intelligence for Engineering Design, Analysis and Manufacturing* 7(2): 87-96.

Fischer, G. 1993. Turning Breakdowns into Opportunities for Creativity. In Proceedings of the International Symposium on Creativity and Cognition, Loughborough, England.

Gero, J. and Maher, M. 1993. *Modeling Creativity and Knowledge-Based Creative Design*. Hillsdale, NJ: Lawrence Erlbaum Associates, Publishers.

Goel, A. and Chandrasekaran, B. 1992. Case-based Design: A Task Analysis. In C. Tong and D. Sriram (eds.), *Artificial Intelligence Approaches to Engineering Design, Volume 2: Innovative Design*. San Diego, CA: Academic Press.

Hinrichs, T. 1992. *Problem Solving in Open Worlds: A Case Study in Design*. Northvale, NJ: Erlbaum.

Johnson, W.L., Benner, K.M., and Harris, D.R. 1993. Developing Formal Specifications From Informal Requirements. *IEEE Expert* 8(4): 82-90.

Jones, E.K. 1992. The Flexible Use of Abstract Knowledge in Planning. Northwestern University, Institute for the Learning Sciences Technical Report no. 28.

Kolodner, J.L. 1993. *Case-Based Reasoning*. San Mateo, CA: Morgan-Kaufman Publishers, Inc.

Kolodner, J. 1983. Reconstructive Memory: A Computer Model. *Cognitive Science* 7(4): 281-328.

Kolodner, J.L. and Penberthy, T.L. 1990. A Case-Based Approach to Creativity in Problem Solving. In Proceedings of the Twelfth Annual Conference of the Cognitive Science Society, Cambridge, MA.

Kolodner, J.L. and Wills, L.M. 1993. Case-Based Creative Design. In AAAI Spring Symposium on AI and Creativity. Stanford, CA. Reprinted in *AISB Quarterly* 85: 50-57.

Koton, P. 1988. Reasoning about evidence in causal explanation. In Proceedings of the 6th National Conference on Artificial Intelligence. Cambridge, MA: AAAI Press/MIT Press.

Lowry, M. 1987. The Abstraction/Implementation Model of Problem Reformulation. In Proceedings of the 10th International Joint Conference on Artificial Intelligence, pp. 1004-1010. Milan, Italy.

Navinchandra, D. 1991. *Exploration and Innovation in Design: Towards a Computational Model*. New York: Springer-Verlag.

Navinchandra, D. 1992. Innovative Design Systems: Where are we, and where do we go from here?. Parts I and II. *The Knowledge Engineering Review* 7(3): 183-213 and 7(4): 345-362.

Prabhakar, S. and Goel, A. 1992. Performance-Driven Creativity in Design: Constraint Discovery, Model Revision, and Case Composition. In Proceedings of the Second International Conference on Computational Models of Creative Design. Heron Island, Australia.

Reubenstein, H.B. and Waters, R.C. 1991. The Requirements Apprentice: Automated Assistance for Requirements Acquisition. *IEEE Transactions on Software Engineering* 17(3): 226-240.

Schank, R. 1982. *Dynamic Memory: A Theory of Learning in Computers and People*. New York: Cambridge University Press.

Seifert, C., Meyer, D., Davidson, N., Patalano, A., and Yaniv, I. 1994. Demystification of Cognitive Insight: Opportunistic Assimilation and the Prepared-Mind Perspective. In R.J. Sternberg and J.E. Davidson (eds.), *The Nature of Insight*. Cambridge, MA: MIT Press. Forthcoming.

Turner, S.R. 1994. *MINSTREL*, Lawrence-Erlbaum Associates, Inc. Forthcoming.