

Indexing Stories as Social Advice *

Eric A. Domeshek
Institute for the Learning Sciences
Northwestern University
Evanston, Illinois 60201

Abstract

This paper reports on an indexing system supporting retrieval of past cases as advice about everyday social problems; it has been implemented in the *Abby* lovelorn advising system. Two points are emphasized: (1) indices are descriptions of problems and their causes, couched in a vocabulary centered on intentional causality, and (2) indices fit a fixed format that allows reification of identity and thematic relationships as features. *Abby* answers several of the central questions that any indexing system must address, and has advantages over less restrictive systems.

Reminding in Everyday Affairs

Much of everyday conversation involves trading stories. This is not always idle chit-chat. Stories usually serve some purpose given the context in which they are told. Often a story can serve as *advice*: it can suggest a course of action, or point out pitfalls to be avoided. Advice, however, need not be explicit, nor need we assume that a story was chosen and told to communicate any particular advice. Instead, advice may be in the mind of the beholder, and interestingly, the *teller* may be one of the beholders. That is, the story may make a point but it need not have been retrieved from memory on the basis of a point to be made.

Consider the following sample reminding. A teenage girl complains: *“My parents won’t let me marry Johnny; they think I’m too young and shouldn’t marry yet.”* It would seem natural for her listener to recall the following story and tell it to her in response:

Sometimes parental concern can be overwhelming. I knew a woman of about 30 who finally had to move out of her parent’s house because her mother kept pestering her to get married. She liked men

*Thanks to Andy Fano, Larry Birnbaum, Eric Jones, and the anonymous AAAI reviewers for many useful comments. This research was supported in part by the Defense Advanced Research Projects Agency, monitored by the Air Force Office of Scientific Research under contract F49620-88-C-0058, and by the Air Force Office of Scientific Research under contract AFOSR-89-0493. The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting, part of The Arthur Andersen Worldwide Organization. The Institute receives additional funding from Ameritech (an Institute Partner) and IBM.

well enough and intended to get married some day. She just hadn’t met anyone she wanted to marry yet (and who wanted to marry her). She had lots of friends, and she went out on dates fairly often; this was no social misfit or recluse. But whenever she got home, her mother was waiting up for her to ask if she had met anyone who interested her. It finally just got to be too much for her, so she had to find a place of her own. While that was probably an overdue move for her anyway, the mother could have avoided forcing the issue if she wanted to keep her daughter around.

This story was chosen by a program named *Abby*, designed to give advice on lovelorn problems. *Abby* is intended to explore how stories surface from memory and assumes that such reminders play a functional role in situation assessment. Recalling a story contributes towards forming a response to a situation; it does not simply yield a cute way to communicate an already derived conclusion. As this example suggests, *Abby* can often find something appropriate to say, even though it lacks processes of evaluation, self-censorship or tailoring of story renditions, which contribute to the quality of people’s everyday advice.¹

Though *Abby* does not choose stories to make a particular point, a little thought reveals that this story does have one to make. Addressed to the girl, the simplest interpretation is that if she dislikes like her parents’ interference in her love-life then she should move out. But here an important mismatch makes itself felt: this is a teenager not a 30-year-old. A reasonable conclusion, to which the girl might be led, is that if she is not old enough to move out on her own then she may not be old enough to ignore her parents’ wishes.

This story was not chosen to illustrate the point *“wait until you are older and can take care of yourself.”* Instead, it was retrieved on the basis of situational similarity under some descriptions of the two situations. Our concern in this paper is to outline constraints on the sorts of descriptions that make good retrieval cues for social advice.

¹Note that the slant of this story rendition — its emphasis on the unfounded nature of the mother’s concern and the focus in the last sentence on the mother’s actions — is an artifact of *Abby*’s use of canned text and its lack of sensitivity to which agent is being advised.

The Indexing Problem(s)

Telling advice stories in response to problematic situations is an instance of case-based reasoning (CBR). Abby's conversational approach emphasizes the process of *case retrieval* over the other sub-processes that compose a complete CBR model [Kolodner et al., 1985, Hammond, 1986]. Basing retrieval on situation descriptions rather than solution descriptions emphasizes how retrieval can contribute to problem solving. Even with the relatively broad latitude afforded by a human interpreter, in any given situation, some stories in memory are going to be more useful than others. One of the unavoidable important issues in CBR, then, is how to arrange that past cases are recalled in appropriate circumstances. This has been called the *indexing problem*.

The basic indexing model of retrieval is very simple. An index *label* assigned to a case in memory is some combination of situational features which together may indicate that the case is worth recalling. An index *probe* is some description of a current situation that is intended to be matched against the labels available in memory. In successful indexing, a match (or partial match) between a current probe and some label indicates that the case associated with the label will offer useful guidance in coping with the current situation.

The appeal of treating case retrieval as an indexing problem is that it provides a simple, understandable model that focuses attention on a series of sub-problems: index content, structure, matching, organization, search, and generation. I have primarily been interested in the questions of content and structure. Indices are defined functionally as the basis for retrieval. Specifying content sufficient to support this function is the next logical step in building an indexing theory. Commitment to specific content implies some minimal structural commitments, but questions of matching, organization, search and generation can largely be deferred. As argued in [Schank et al., 1990], defining index content is the largest task in constructing an indexing system yet it has received relatively little attention in the literature. For the social domain and the task of advising on problems, the notable exception is the work of Seifert [Seifert, 1987].

At the most general level, prior research in CBR suggests that to find cases that can help accomplish a task, the most effective indices are descriptions of task-related goals and features of the world causally relevant to the status of those goals [Hammond, 1986]. Hammond's cooking program indexes cases on desired properties of dishes and features of ingredients that affect those properties. A lovelorn advisor ought to index cases on social goals that can run into difficulties and on features of social situations that lead to such difficulties. The task before us is to move from this general claim about functional indices towards specific answers to the first questions any indexing theory must address: what is the allowable content of a set of indices and the actual content of particular indices.

Abby's Indices

As a first pass, we can say that Abby labels cases as follows: a label describes the general problem, (mother invades daughter's privacy), the flip side of the problem, (mother wants to prompt daughter to marry), and some causally relevant background, (the daughter's age accounts for the mother's insistence). Such labels meet the twin criteria of availability and usefulness. When understanding a new situation, the available or derivable features from which a probe might be constructed will typically include such components — the goals subject to positive and negative impacts plus some causal antecedents of the pivotal actions. When matching probes and labels, commonalities on such features strongly indicate the relevance of the label's story; the story can advance problem solving and finding it does not depend upon some prior solution.

Intentional Causality

Most events in the everyday world happen at the behest of one agent or another. Agents — individuals, groups and institutions — take action because they expect and desire particular outcomes. Index descriptions must include intentional causation if they are to capture the most important mechanisms at work in these everyday social situations. We require a vocabulary for intentional causes and, ideally, a canonical format for organizing elements of the vocabulary. Fortunately, earlier research on understanding stories of everyday events developed both a vocabulary for intentional explanations of agents' actions [Schank and Abelson, 1977] and a standard format for such explanations called the *intentional chain* [Schank, 1986].

The intentional chain is designed to explain why an action was performed, or how a state resulting from an action came to be. Such a chain accounts for an action as a step in some plan designed to accomplish a particular goal which is attributed to an agent on the basis some longer term personal or relational theme. Consider the complaint of the teenage girl from the opening example; we might account for the girl's *not being married* (EFFECT) as a result of her parents *forbidding her to marry* (ACTION) which was their chosen method of *obtaining her compliance* (PLAN) with their *wish that she not be married* (GOAL) stemming from their responsibility towards her in *their role as parents* (THEME). Figure 1 shows, side by side, the template for an intentional chain and how that structure might represent this explanation of the opening problem.

The chain pictured on the right in figure 1 is part of the probe for our initial problem situation; if a story tells how similar motivations leading to a similar problem were once overcome, we want to hear that story. Such a story might tell us what to do when parents forbid children to do something, when parents take a stand on their children marrying, or when parents create problems by acting on their parental responsibilities. In fact, the story Abby retrieved told of a parent *prodding* (not forbidding) her child to marry, but other motivational similarities mitigated this difference.

Slots	Contents
<i>Theme</i>	theme
<i>Goal</i>	goal
<i>Plan</i>	plan
<i>Action</i>	action
<i>Effect</i>	state

Slots	Contents
<i>Theme</i>	parent-of
<i>Goal</i>	P:not-wife-of
<i>Plan</i>	threaten
<i>Action</i>	forbid-wife-of
<i>Effect</i>	not-wife-of

Figure 1: Generic Intentional Chain and Example

This chain is not, however, enough to stand on its own as an index. It is part of a causal account of how the problem came to be, but it does not include a description of the problem itself. A *problem*, in the everyday world, is some actual or potential negative effect on a goal. The pictured chain only talks about the successful pursuit of the goal to keep the daughter unmarried, while the problem is the failure of the daughter's goal to be married. Since this chain only talks about why the parent forbade the marriage, an index including just this chain provides no basis to prefer a story about parental action interfering with a child's desires over, say, a story where one parent's action offended the other parent. Therefore, in addition to the pictured chain, the index must include the problematic impact of the effect: that the girl's goal to be married was frustrated. For **Abby**, with its task of giving advice, we will always want our indices to include problems since the problem is what calls for advice.

Including a problem in an index can be accomplished by including something very like another intentional chain — the negatively impacted goal and the state that affected it, both ideally accompanied by their causal antecedents (a theme and an action respectively). But these components together form a kind of “mutant” intentional chain. That is, though such representational groupings normally explain an action or state, these groupings somehow lack the causal coherence to form a connected chain; the gap in the chain indicates the existence of a problem. In this example, we can construct the following package: the girl's *relationship to her boyfriend* (THEME) led to her *wanting to marry him* (GOAL) but her parents *forbidding her to marry* (ACTION) led to her *not being married* (EFFECT).

One advantage of packaging our problem description this way is that we can pair it up with the original intentional chain to picture the combination as a *chain-interaction*. Figure 2 shows what this looks like. The basic idea is to show how a single action can have multiple impacts when considered with respect to different goals (or different effects); the defining feature of a chain-interaction is that two intentional chains (possibly mutant chains) hinge on the same action (thus, in figure 2, the single ACTION should be understood as belonging equally to both columns)². This is a way of elaborating and displaying something like Wilensky's idea of a goal interaction [Wilensky, 1979]. As this ex-

²Also, note that the lack of a plan in the second column is due to our ignorance about how the teenager was pursuing her goal when her parents actions messed things up for her.

	Chain-1	Chain-2
<i>Theme</i>	parent-of	<i>Theme</i> lover-of
<i>Goal</i>	P:not-wife-of	<i>Goal</i> A:wife-of
<i>Plan</i>	threaten	<i>Plan</i>
<i>Action</i>	forbid-wife-of	
<i>Effect</i>	not-wife-of	<i>Effect</i> not-wife-of

Figure 2: Chain-Interaction for Input Situation

ample suggests, chain-interactions make good indices because they package a problem with an account of why the problem was produced. As noted, **Abby**'s indices will always include problems and problems usually result from motivated actions; we should consider the positive when trying to eliminate the negative.

Often a problem has a history or background beyond the immediate motivation and consequences of a single pivotal action. In our opening example, the girl's age counts as a relevant piece of background information because it helps account for her parents' actions. Her age and its causal relationship to the problem ought to be part of the index. Relevant background can, however, be far more complicated. Consider a man bemoaning the fact that his wife is leaving him; if the man has been carrying on an affair, the affair is background that ought to affect a choice of advice, thus the affair and its causal influence on the ultimate divorce must appear in the index. Here, the causal connection is that the affair was problematic for the wife — that it was a violation of their relationship which led her to dissolve the marriage. This exemplifies one common sort of background: a provocation which leads to a problematic response. But note that such a provocation can usefully be described as a problem in its own right, and further, that it can usefully be elaborated by including the positive motivation for the problematic act, (why the husband had the affair). The provocation, then, can be described by another chain-interaction: *the man wanted sexual satisfaction and had an affair to get it; meanwhile the woman wanted loyalty from her husband but his affair violated this obligation*.

This example shows why an index might usefully comprise two chain-interactions, the first built around an action causally related to the problematic action appearing in the second. It seems, though, that we could endlessly invent situations hinging on ever more distant causal antecedents. Why not argue that an index might usefully include the fact that the man chose to have an affair because his father had before him, and that his father had affairs because his wife ignored him, because she did not respect him, because he lost his job, because... There is an increasing cost and diminishing return for such index expansionism. It gets progressively harder to code such indices by hand or imagine how a system might generate, organize, and match them. Meanwhile, given that case retrieval need only choose from a closed set and need not be perfect, the additional precision may not be necessary or particularly helpful. Given the type and number of cases **Abby** must discriminate, two linked chain-interactions appears to be as complicated as an index need get.

Slots	Contents	Slots	Contents
<i>Theme</i>	theme	<i>Theme</i>	parent-of
<i>Holder</i>	agent	<i>Holder</i>	adult
<i>Partner</i>	agent	<i>Partner</i>	teenage-girl
<i>Thread</i>	thread	<i>Thread</i>	family-life
<i>Goal</i>	goal	<i>Goal</i>	P:not-wife-of
<i>Plan</i>	plan	<i>Plan</i>	threaten
<i>Action</i>	action	<i>Action</i>	forbid-wife-of
<i>Effect</i>	state	<i>Effect</i>	not-wife-of

Figure 3: Reprise of Figure 1 with Expanded THEME

Setting Limits: A Fixed Index Structure

What we are aiming for is a clear statement of the content and structure of indices. That includes a limit — an upper bound — on index complexity. The structural limit of two linked chain-interactions is a good start, but we also must say what can fill the THEME, GOAL, PLAN, ACTION, and EFFECT slots. Using standard composable representations and graph matching techniques, the information accessed through an index could easily remain unbounded.

My approach has been to close this trap door by deciding up front what amount of substructure may be used in describing each of the basic chain components and explicitly including fields for just that substructure in the index. Figure 3, for example, replicates figure 1 but shows the expansion of the THEME component of the intentional chain. Here, the THEME field holds a characterization of the nature of the theme, the HOLDER contains the agent whose theme it is, the PARTNER contains the agent related to the holder by the theme, and the THREAD holds a characterization of the long-term pattern to which this theme belongs.

Once this expansion is completed, the filler of each index field can be treated essentially as an atom. For purposes of comparing fillers however, we borrow the notion of microfeatures [Hinton, 1981], encoding the filler’s type as a pattern from an appropriate closed set of discriminations. Agents, for example, are patterns that can discriminate sex, age, number, etc. These encodings however are the absolute bottoming out of type information in an index. **Abby**’s indices cannot for instance mention how long an agent held a theme, or identify some prior theme from which it evolved.

There is however another form of information encoded in indices: *relational information*. Some index field fillers are interpretable as tokens that can appear in more than one slot. For example, the same goal may appear more than once in an index, perhaps as the goal suffering in the first chain-interaction and then as the goal benefiting in the second; this is a common version of the provocation and response pattern — *damage and restoration*. Based on this pattern, we might find a story about giving the sufferer an alternate compensation for leaving the initial problem uncorrected. For retrieval to be sensitive to such patterns the relationship features must be part of the indices. **Abby** reifies all such binding relationships making them explicit features; the strict limit on fields makes it possible to

enumerate all the possible relationships of this type.

Agents are important tokens in the social world, and are related in particularly interesting ways. Rather than simply noticing whether one agent is the same as another or is someone else, **Abby**’s indices record thematic characterizations of their relationships. For instance, our probe index records that the agent performing the action is the parent of the agent whose goal suffers from that action. In matching a probe with this pattern, **Abby** will prefer labels where the actor and suffering agent are related in a similar way. Similarity here is defined just as for the fillers of the THEME fields — by amount of overlap in the microfeature encodings of the theme characterizations.

By virtue of their strict use of atomic slot-fillers, fixed structure, and thus fixed relationship possibilities **Abby**’s indices can be viewed as semantically rich binary feature-vectors. Similarity judgement can be reduced to microfeature overlap scores. **Abby** uses a formula that is a variant of the similarity metric proposed in [Tversky, 1988]: the score for each label is calculated as $2 \times |probe \cap label| - |probe - label| - |label - probe|$. Simply combining the counts of the feature overlap and difference sets sidesteps the problem of fiddling with individual feature weights; the only violation to this simplicity in the current implementation is that fillers of GOAL states receive a constant multiplicative boost.

Abby retrieves cases simply on the basis of these similarity scores. It does not implement any filtering or adaptation processes. Nonetheless, the stories it tells can often be interpreted as offering relevant advice.

Sample Indices

This paper cannot present and justify the detailed structure of **Abby**’s indices. Nonetheless, figure 4, outlining the probe encoding the teenage girl’s problem and the label attached to the retrieved story is included to complete the sketch of the example reminding. The greatest difference between this probe and label turns out to be the specific problems described. The probe is about the failure to marry while the label is about a lack of privacy. Nonetheless, the influences leading to these problems are sufficiently similar that this difference was overcome both in **Abby**’s retrieval and in our interpretation. In other situations, the problems might be more similar and the causality less so. Ideally, a label would exist in memory that matched on both the problem and its antecedents.

In addition to the obvious similarity in their patterns of slot fillers, these indices share many filler relationships. For instance, following a typical pattern, in the left column of both indices’ bottom chain-interaction, the agent who acts also does the planning, holds the goal being served, and holds the theme from which the goal stems. In both indices, the actor’s theme relates to an agent who is the object of the goal, the recipient of the action and the object of the effect. This related agent is also the holder of the suffering theme and goal in the right column, as well as the subject of the background and problematic effects. Together with

Probe Index for Teenager's Problem

<i>Effect</i> age-low	<i>Effect</i>
<i>Link</i> ↖-effect-bias-/-plan	
<i>Theme</i> parent-of	<i>Theme</i> lover-of
<i>Goal</i> P:not-wife-of	<i>Goal</i> A:wife-of
<i>Plan</i> threaten	<i>Plan</i>
<i>Action</i> forbid-wife-of	
<i>Effect</i> not-wife-of	<i>Effect</i> /-effect

Label Index for Nosy Mother

<i>Effect</i> age-medium	<i>Effect</i>
<i>Link</i> ↖effect-bias-/-plan	
<i>Theme</i> mother-of	<i>Theme</i> be-young-woman
<i>Goal</i> A:wife-of	<i>Goal</i> S:privacy
<i>Plan</i> prod-partner	<i>Plan</i>
<i>Action</i> request-lover-of	
<i>Effect</i> lover-of	<i>Effect</i> privacy-low

Figure 4: Sketch of Example Probe and Label the common parent-of theme, these binding similarities represent the commonality: *acting in their parental role, a parent does something to a child on account of the child's age, that interferes with the child's goals.* This degree of match helps us start extracting relevant advice from the reminding, though as suggested earlier, the age differences would likely end up reshaping the ultimate recommendation.

It is instructive at this point to consider some possible variants of these indices and look at how the representable distinctions could contribute to retrieving relevant advice.

- Imagine that **Abby** had another story with an index just like the retrieved label except that the badgered child was a teenager just like the sufferer in our probe. Such a label would have surfaced as a better match, and we would expect its story to be more directly applicable since it would respect the constraints on what a teenager can reasonably do in opposition to parental pressure. It might *directly* suggest waiting until they relented or until they could be openly defied, or it might suggest playing on their love, for example, by pining for the forbidden lover.
- Imagine **Abby** had a label more like the probe in that the child's goal directly conflicted with the parents', was rooted in a relationship to a third party, and was held mutually with that other theme partner. If such a label existed and won out its story might suggest ways in which the child's partner could help convince the parents to drop their opposition.³
- Consider a probe situation where the goal at stake was the establishment of the child in a job rather than a marriage; since **Abby** explicitly represents

³Note that the described overlap, easily representable in **Abby's** indices, fits the thematic pattern **mutual-goal / outside-opposition** suggested in [Schank, 1982] as a possible support for cross-contextual reminders.

many features underlying a wide range of relationships, including, for instance, that people take and leave jobs with more ease than marital relations, it might now find a story suggesting that the child be allowed to take the job on a trial basis.

Discussion

In discussing variants of the example indices just now, my intent was to show how some of the particular distinctions expressible in these fields, (the age of the suffering agent; the characterization of goals as mutual; the binding relationships revealing outside opposition; relationships' differing barriers to entry and exit), are useful for discriminating relevant advice. I made similar arguments as I introduced the overall shape of **Abby's** index frame, to justify inclusion of the frustrated goal, explanations for that goal's origin and for the problematic action, and some selected background.

Such arguments exemplify a general methodology for justifying representational content: construct cases that differ on some feature and show that a system's behavior ought to be sensitive to the distinction. This approach also suggests a data-driven methodology for *developing* representations in the first place: look at many cases and determine what features account for their similarities and differences. **Abby** illustrates a case-based approach to building a case-based reasoner.

The feature-by-feature, case-by-case style of argumentation is ill suited to short presentations, such as this current paper. Given a complete system it is sometimes desirable and convenient to use a different, more coarse-grained, style of argument: an empirical evaluation. At this point, however, I can only offer the most crude and preliminary study, conducted informally at a time when the system contained 100 labels. Presented with 5 probes, 80% of **Abby's** top rated stories (considering for each probe the top three scoring stories) were judged acceptable advice for the problems those probes encoded. Acceptability was determined by rating the quality of advice extracted from a story, given an English statement of the probe's problem and one minute to read and think about the story as advice for that problem. Applying the same criteria to stories sampled at random suggested that, on average, only 20% of the corpus would have been acceptable. Not much is claimed for this study other than that it is encouraging. **Abby** currently contains about 250 labels and is due for more extensive and systematic evaluation when it reaches its immediate target of 500 labels.

Abby offers a detailed theory of index content and structure. Though indices are simply flat feature vectors, they can be clearly described as organized compounds of recognized representational units. The description in terms of intentional chains and their components begins the commitment to specific allowable content, a commitment carried through in much greater detail than this paper can reflect. That content will, in the end, be among the more important contributions of the **Abby** project. Establishing a set of distinctions that capture much of what matters in the

social world is a difficult and important task. Yet this area has received relatively little attention outside the Natural Language Processing community, where early concern with texts describing everyday life forced the issue. In contrast, major sub-industries have been established in AI to work out appropriate representations for reasoning about the physical world, and to some extent the mental world [Hobbs and Moore, 1985].

Abby's use of indexing to drive development of representations for the social world is most closely paralleled by the proposal in [Schank et al., 1990] for a *universal index frame* (UIF). This is unsurprising, since I was involved in the design of the UIF. Abby's index frame can be viewed as a variant of the UIF extended to allow two-step temporal sequences, but restricted to consider only two-way comparisons built around common actions. Details of fields and fillers also differ somewhat; Abby neglects parts of the UIF while developing others in more detail. Much of Abby's index vocabulary predates the design of the UIF.

Many of these systems' differences stem from the UIF's commitment to universality and Abby's need to accomplish its particular task. The UIF emphasized the notion of *anomaly* (rooted in CBR's concern with learning) while Abby focused on the more restricted concept of *problems*. The UIF design effort carried its focus on content to the point of intentionally avoiding consideration of processes that might operate on the structure; since it seemed there were examples where comparisons between three or more versions of the same event were salient, the UIF placed no constraint on the number of parallel intentional chains. With Abby, on the other hand, I decided to explicitly accept limitations, both for the sake of clarity and because it allowed an efficient matching algorithm.

Abby's indices promote clarity by explicitly committing to a set of features that can affect retrieval; the only influences on reminding are the microfeature definitions of the allowed fillers in their designated fields, and the reified relational features encoding multiple references to the same token or social relationships between agents. Nothing is hidden in a complicated matching procedure traversing unbounded recursive structure.

Given this index structure, the matching process can be implemented in a very simple, massively parallel, "connectionist" style system. Abby is a demonstration of how semantically rich domain knowledge can be encoded into a feature vector format without missing the meat of the domain; it suggests how an interesting task could successfully be handled by a connectionist system. In the long run, connectionist implementations are attractive because they offer the potential advantages of scalable performance, a simple scheme for handling partial matching, a clearer correspondence to biological systems, and hope for performance improvement through feedback-driven weight tuning.

Feature vectors also support economy in the important process of *probe generation*. The ANON system [Owens, 1988] showed how massive parallelism could

help control inference of features required to discriminate among a set of possible retrieval targets. ANON approximated an optimal discrimination network no matter what order information became available; it could always choose to direct inference towards the establishment of features present in approximately half the remaining contenders. One requirement of the algorithm was that all inferable features could be referenced unambiguously from any label, enabling the system to determine how many labels contained any feature. A feature vector format meets this condition trivially.

Abby illustrates an approach to indexing that is useful and effective, that yields specific answers to the most basic questions about index content and structure, that suggests a (potentially) efficient implementation of matching, and that dovetails nicely with studies of probe generation. Time and experience will reveal this systems' ultimate retrieval accuracy; time and new computer hardware will determine its practicality.

References

- [Hammond, 1986] Hammond, K. (1986). *Case-based Planning*. PhD thesis, Yale University.
- [Hinton, 1981] Hinton, G. (1981). Implementing semantic networks in parallel hardware. In G.E., H. and Anderson, J., editors, *Parallel Models of Associative Memory*, chapter 6, pages 161-187. Earlbaum, Hillsdale, NJ.
- [Hobbs and Moore, 1985] Hobbs, J. and Moore, R. (1985). *Formal Theories of the Commonsense World*. Ablex, Norwood, N.J.
- [Kolodner et al., 1985] Kolodner, J., Simpson, R., and Sycara-Cyranski, K. (1985). A process model of case-based reasoning in problem-solving. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA. IJCAI.
- [Owens, 1988] Owens, C. (1988). Domain-independent prototype cases for planning. In *Proceedings of the First Workshop on Case based Reasoning*, Clearwater, FL. DARPA.
- [Schank, 1982] Schank, R. (1982). *Dynamic Memory*. Cambridge University Press.
- [Schank, 1986] Schank, R. (1986). *Explanation Patterns*. Earlbaum, Hillsdale, NJ.
- [Schank and Abelson, 1977] Schank, R. and Abelson, R. (1977). *Scripts, Plans, Goals, and Understanding*. Earlbaum, Hillsdale, N.J.
- [Schank et al., 1990] Schank, R., Brand, M., Burke, R., Domeshek, E., Edelson, D., Ferguson, W., Freed, M., Jona, M., Krulwich, B., Ohmaye, E., Osgood, R., and Pryor, L. (1990). Towards a general content theory of indices. In *Proceedings of the AAAI Spring Symposium on Case-Based Reasoning*, Palo Alto, CA. AAAI.
- [Seifert, 1987] Seifert, C. (1987). *Mental Representations of Social Knowledge*. PhD thesis, Yale University.
- [Tversky, 1988] Tversky, A. (1988). Features of similarity. In Collins, A. and Smith, E., editors, *Readings in Cognitive Science*, pages 290-302. Morgan Kaufmann, San Mateo, CA.
- [Wilensky, 1979] Wilensky, R. (1979). *Understanding Goal-Based Stories*. PhD thesis, Yale University.