Using Points to Construct Browsing Links in ASK Systems

Chip Cleary and Ray Bareiss
The Institute for the Learning Sciences
Northwestern University
1890 Maple Avenue
Evanston, IL 60201

chip@ils.nwu.edu, bareiss@ils.nwu.edu

Abstract

Our research concerns the construction of knowledge-rich memories, in hypermedia form, for use as aids to problem-solving. One of the most difficult steps in building such memories is constructing a rich set of links between the content elements they contain (i.e., text segments, graphics, and video clips). This paper describes point linking, a method we have developed for automated linking. Point linking is currently being incorporated into the ASKTool, a hypermedia editor in ongoing use at the Institute for the Learning Sciences.

1. Introduction

The field of AI has had limited success in building traditional knowledge-based systems. Such systems typically require comprehensive, explicit representations of both the domain in which the system is to operate and the task which the system is to perform. Unfortunately, constructing such representations is often beyond the current state of the art. A more pragmatic approach towards building pragmatically useful knowledge-based systems is to construct computerized memories which contain expertise that a human problem solver can tap into when he or she needs information or advice. Our research is concerned with the construction of ASK systems, a type of hypermedia-based computerized memory for use as aids to problem solving.

In this paper, we describe some of the progress we have made over the past three years in formalizing and automating the process of building links between stories[1] in ASK systems. We first describe and evaluate narrative linking, an initial method for automating linking which used a complex AI-inspired representational framework. Although narrative linking performed well, it was difficult to use. We then discuss point linking, a method we developed with the design objective of using "just enough" AI to produce a successful linker. Point linking is currently being incorporated into the ASKTool, a hypermedia editor in ongoing use at the Institute for the Learning Sciences (ILS).

2. The ASK Approach to Browsing

In contrast to most hypermedia systems, ASK systems are based on the metaphor of a conversation with an expert (Bareiss & Osgood, 1993). In particular, they are designed to simulate a question-answer dialog in which the user asks the questions and the system provides the answers. An interaction with an ASK system consists of two phases: first a user zooms to an initial story relevant to his interests, then he browses the database by asking follow-up questions as desired and retrieving additional stories in response. If the user wants to redirect the "conversation," he may return to the top-level of the system to begin another round of zooming and browsing.

To support browsing, ASK systems contain a rich network of links, each of which joins a source story which raises a question to a target story which answers it. The browsing interface of an ASK system surrounds the current story with specific questions it might raise for the user. Eachof these questions mark a browsing link; clicking on a question takes the user to a target story which answers it. The questions are grouped by by generic question type. If a user has a specific question in mind, these types enable him or her to quickly find it. If the user has only a vague idea of what question to ask, he or she can simply scan through the questions under a generic question that looks promising.

The typology of generic questions used in ASK systems are inspired by a simple theory of conversation which argues that at any point in a conversation, there are only a

[1] We use "story" colloquially to refer to a single content element in an ASK system (i.e., a piece of text, segment of video, or graphic).

Refocusing: Adjustments to the specificity of topic under consideration.
1. Context: What is the big picture within which the current topic fits?
2. Specifics: What are the details of this situation or an example of it?
Comparison: Related topics at the same level of abstraction as the current one.
3. Analogies: When have other similar situations occurred?
4. Alternatives: What other approaches exist, or what did other experts say?
Causality: Explanations and outcomes
5. Causes (or earlier events): How did this situation develop?
6. Results (or later events): What is the outcome of this situation?.
Advice: Planning knowledge for use in the problem solver's situation.
7. Opportunities: How can I capitalize on this situation?
8. Warnings: What should I watch for that might go wrong?

Figure 1: The Eight CACs Used in the TransASK System

few general categories of follow-up statements that constitute a natural continuation rather than a topic shift (Schank, 1977). These "conversational associative categories" (CACs) can also be thought of as the general classes of questions a person is likely to formulate in conversation. The particular CACs used in ASK systems are tailored for conversations about problem-solving. The TransASK system, for example, employs the eight CACs shown in Figure 1. Other ASK systems use similar (and often identical) sets.

3. The Goals of Automated Linking

ASK systems are difficult to build, taking trained teams of indexers person-months (and sometimes person-years) to construct. The two most time-consuming and expertise-intensive steps in building an ASK system are acquiringappropriate stories and generating a rich set of links between them. Cleary and Bareiss (1994) describes the knowledge-acquisition process; here we concentrate on the problem of building links.

As a step towards a system which can link stories autonomously, our short-term goal has been to provide an automated linker which operates in partnership with a person to build browsing links between stories. We envision that the person will first construct representations of stories, the linker will then suggest links between them, and finally the person will review and edit those links. Figure 2 uses an example link from the Engines for Education ASK system to illustrate the specific tasks an automated linker must perform (Schank & Cleary, 1995).[2]

Three criteria are critical when evaluating the results produced by an automated linker. *Thoroughness* measures how completely the linker supports the three tasks in the linking process. *Recall rate* measures the percentage of the links which an ASK system should have that the linker actually generates. *Precision rate* measures the percentage of links that a linker actually generates which are links that an ASK system should have. An additional criterion is useful to judge the input a linker requires. *Ease of use* judges how hard is it for an indexer to learn and use the representational scheme a linker uses.[3]

Note that the relative importance of these criteria must be considered in the context of a target audience of indexers. Subject matter experts who build ASK systems as job aids require a linker that is easy to use even if it performs only moderately well on the results criteria. Professional indexers can invest the time to learn and use a more complicated linker if it produces better results.

4. Narrative Linking

Our initial attempt at automating linking, narrative linking, employed a rich representation which treated stories as narratives about planful behavior. This method represented stories using a "narrative frame" (Osgood, 1994; Osgood & Bareiss, 1993) which encapsulated a naive model of intentionality inspired by the "intentional chain" of (Schank & Abelson, 1977). The narrative frame provided a fixed set of domain-independent slots (including, among numerous others, *AgentRole*, *Goal*, *Plan*, *Enablements*, *Impediments*, *Anticipated Actions*, and *Unanticipated Outcomes*).

To represent a story, an indexer instantiated the frame by inserting domain-specific fillers into its slots. To construct links between stories, the automated linker employed "narrative linking rules," each of which could infer links for a specialized sense of one of the eight CACs. Conceptually, a rule performs a pairwise comparison between the frames representing two stories. For example, a rule might specify that two stories whose agents who have the same *Goal* but employ different *Plans* should be linked as *Alternatives*.

We ran an informal experiment to see whether narrative linking would be effective when used by novice indexers. A

2 Engines for Education deals with the role that technology can play in reforming education. It contains approximately 350 stories and 4000 links.

3 When judging recall and precision, we use the "gold standard" of systems built manually by human indexers. Although manually-constructed systems are not perfect, they do provide an objective baseline. Furthermore, manually-constructed systems are not perfect in a predictable way -- they typically include few bad links but miss a number of good ones. In other words, when judged against an idealized ASK system, manually-constructed systems have a high precision rate but only moderate recall rate. These distortions have little impact when evaluating the recall rate of an automated linker. However, they provide an overly conservative measure of precision.

Source Story: Different Simulators for Different Skills
Tools such as the flight simulator pave the way for a natural, effective way to learn physical skills. To teach social skills in a similarly effective way, we need to build social simulators....

CAC: Specifics
Question: "What is an example of a social simulation?"

Target Story: Dustin
We have built a number of social simulations that provide learning-by-doing environments. Dustin is an example of a simple simulator that helps students learn a foreign language through using it....

Linking Task	Output for Example Link
1) Determine that the two stories are related	They are both about social simulation
2) Determine how they are related (i.e., which CAC joins them)	The target story provides *specifics* about the source story
3) Build a question to bridge from the source story to the target story	"What is an example of a social simulation?"

Figure 2: An example link illustrating the three tasks a linker must perform

team of two Northwestern University undergraduates were given a corpus of 47 stories in the domain of military logistics (a subject about which they knew nothing) and the empty narrative frame. They created taxonomies of fillers, instantiated the frame for each story, and defined inference rules for four of the eight CACs. We then ran the rules on the frames, thereby producing approximately 120 story-to-story links.

Narrative linking proved to be highly precise. 72% of the stories it joined were also linked in the manually-linked ASK system from which the stories were drawn. This rate is over double that of chance -- two stories selected at random from the sample set had only a 37% chance of being manually linked. These results indicated that novices could build narrative frames and that an automated linker could use them to construct links with high precision (i.e., 20 incorrect links out of 120 proposed). Further, narrative linking was thorough, supporting each of the three linking tasks: relating stories, assigning conversational categories, and generating questions from a representational template associated with each linking rule.

Although the experiment did not provide a measure for recall, narrative linking's recall rate will be limited because narrative frames cannot conveniently encode information which does not revolve around intentionality. Nevertheless, the primary problem with narrative linking is that the method is not easy to use. The students worked approximately a total of 160 person-hours to represent 40 stories. An experienced indexer could link the stories by hand in roughly half the time it took to the subjects to engineer the representation!

5. Point Linking

Point linking is the result of our efforts to develop a linking method for professional indexers which, like narrative linking, uses structured representations to achieve high precision and thoroughness but which is easier to use and can handle any type of story, including those that do not involve intentional behavior.[4] Point linking suggests links between stories by inferring question-answer relationships between the points attached to those stories.

One source of motivation for point linking comes from our observation of how experienced indexers behave when linking stories manually. When indexers argue about how two stories should be linked, they often frame the argument in terms of terse statements about the contents of those stories. An indexer might say, for instance, "We should add an *Alternatives* link because the first story states that 'necessity leads to innovation' and the second states that 'hard work leads to innovation.'" Statements like "necessity leads to invention" are exactly what we aim to capture in a representation of points. Why not ground a language used to support linking in the sorts of explanations that expert indexers use?

Another source of motivation comes from observing how good readers behave. When good readers read a text, they boil it down, abstracting its main points.[5] These points help readers to determine what questions they should ask themselves about a text and to integrate what they read with what they already know. If people use points internally to relate new knowledge to pre-existing knowledge, why not use them in an automated system which does the same?

5.1. The Architecture of a Point

We define a "point" of a story as a central piece of information someone might try to convey by telling the

4 See Cleary & Bareiss, (in preparation) for a description of two automated linking methods we have developed for novice linkers. These methods do not captured structured relationships between the concepts in stories; rather they focus on how stories treat concepts individually. Although less powerful than point linking, these methods are easier to learn and require less representational effort from indexers than the methods described in this paper.

5 Collins, Brown, & Newman (1983) stresses the importance of teaching novice readers to build summaries when reading.

Slot Name	Legal Fillers	Example
Concept-1	(Any concept)	Simulations
Mode	Do, Should, Can	Do
Sense	Indeed, Not, Anti	Indeed
Relation	(Any predefined relation)	Enable
Concept-2	(Any concept)	Learning-By-Doing

Table 1. The point frame

story. To illustrate, some example points from Engines are "Simulations enable students to learn-by-doing" and "Multiple-choice-testing is bad." Others have defined "point" differently. For example, Schank, Collins, Davis, Johnson, Lytinen, & Reiser (1982). proposed that the point of a statement was the impact that the speaker intended the statement to have on the listener.

To represent points, we have developed a simple, fixed frame which contains five slots (Table 1). The goal of this frame is not to capture every nuance of the points authors make with stories. Rather it is to capture a limited amount of information (i.e., an amount which is not onerous for indexers to represent) which captures the broad sense of the author's intentions in a form which enables an automated linker to create browsing links.The two concept slots indicate what topics an author addresses in a point. To provide guidance, the representation language for points provides a predefined concept hierarchy (Table 2) which indexers may use as general categories under which they may install specialized domain-specific topics. However, indexers are expected to expand the set of legal concepts to reflect the universe of discourse in their particular ASK system. Furthermore, the goal of the topic hierarchy is not to include every possible type of concept, but rather to provide a skeletal framework which contains just the types of concepts that are most important in linking.

We have found two types of concepts to be particularly useful for linking: goal/plans and agents (which is not surprising since the CACs in ASK systems are designed to support problem-solving). Hence, those sections of the predefined hierarchy are richer than others.

The mode, sense, and relation fields indicate what an author says about the topics in a point (i.e., how the concepts in a point relate to each other). We claim that there are only a small number of important ways in which speakers (and authors) relate concepts to each other. So, the language provides comprehensive, fixed vocabularies for the three fields which deal with how concepts interrelate. The example point demonstrates the strength of this approach. It seems reasonable that a language should contain predefined terms for *Do*, *Indeed*, and *Enable*, but unreasonable to expect it to include *Simulations* and *Learning-By-Doing*.

The mode and sense fields allow an indexer to twist the meaning of a point. The mode field corresponds to logical modality. ,*Does* indicates that the point is known to be true, *Can*that it could possibly be true, and *Should*that it should be made to be true. The sense field indicates the truth value of the point. *Indeed*, the "null" sense, indicates that the point is true, *Not* that it is false, and *Anti* that the inverse of the point is true. *Anti* is included to reduce the number of relations the point language requires. For example, the point *Memorization Does Anti Cause Learning* would mean the same as *Memorization Does Indeed Prevent Learning*. However, *Prevent* is not a legal relation. Because the point language contains *Anti*, it does not need to contain "inverse" relations such as *Prevent*.

The relation field provides the pivot around which the meaning may be twisted. The point language provides a predefined hierarchy of relations (Table 3). The relations it contains are intended to capture (sometimes at a high level) the important conceptual relationships which hold between concepts in the points people make. Unlike the concept hierarchy, this one is not intended to be routinely extended by users. We have encoded the stories in Engines using the point language and have found the current hierarchy

Attribute	Object
Mental Attribute	Conceptual Object
Domain	Policy
Event	Role
Mental Event	Theory
Goal/Plan	Tool
Acquiring Something	Value Judgment
Building Something	Physical Object
Changing Something	Agent
Managing Something	Organization
Measuring Something	Formal Organization
Preventing Something	Informal Organization
Providing Something	Person
Pursuing Something	State
Parable	

Table 2: The Predefined Hierarchy of Concepts

Have-Causal-Relation	Have-Interpretation-Relation	Have-Storytelling-Relation
Enable	Compare-With	Have-Coda
Explain	Is-Better-Than	Have-Description
Have-Result	Have-Value-Judgment	Have-Definition
Implement	Have-Planning-Relation	Have-Illustration
Motivate	Have-Goal	Have-Teaser
Explain	Have-Function	Have-Time-Relation
Have-Class-Relation	Have-Opportunity	Have-History
Have-Classification-Element	Have-Warning	Have-Prognosis
Have-Example	Have-Limitation	Precede
Have-Subclass	Have-Resolved-Warning	
Have-Detail	Use-Plan	
Contain	Apply-Theory	
Have-Attribute	Contain-Step	
Have-Context	Have-Script	
Have-Economic-Relation	Use-As-Policy	
Create	Use-Role	
Possess	Use-Thing	
Know	Borrow	
	Consume	

Table 3: The Predefined Hierarchy of Relations

If:	Source Point	=	[<Concept-1> Does Indeed Contain <Concept-2>]
	Target Point	=	[<Concept-3> Does Indeed Contain <Concept-2>]
Then:	CAC	=>	Alternatives
	Question	=>	"What else also contains a <Concept-2>?"

Figure 3: An example linking rule

sufficient to capture the relationships required for the approximately 750 points which resulted.

5.2. Using Points to Build Browsing Links

A browsing link joins a source story which raises a question to a target story which answers it. To create a new link, point linking applies this progression not at the level of entire stories but rather at the level of individual points -- it joins source points which raise questions to target points which answer them. As an example, consider the "Dustin" story (Figure 2), which makes the point *Dustin Does Indeed Contain Social Simulation*. This source point raises a variety of questions, one of which is the *alternatives* question "What else also contains a social simulation?" This question may be answered by a range of target points, specifically those that match the pattern *Topic-X Does Indeed Contain Social Simulation*. After generating the question, the linker will search for target points whch match the pattern. Another story in the system, the "George" story, contains the point *George Does Indeed Contain Social Simulation.* Since this point matches the pattern and answers the question, the linker draws a link between the "Dustin" and "George" stories, labeling it with the appropriate question and CAC.

Indexers create the points required by this procedure, but where do the questions come from? We are currently constructing a library of point linking rules which formalize these knowledge-intensive questions. The rule in Figure 3 captures the question used in the above example.

We do not yet have empirical results from point linking. However, we expect that point linking will have a precision rate which approaches that of narrative linking because both methods use structured descriptions and explicit inference rules to isolate the most important types of links between stories. We also expect that it will have a recall rate above that of narrative linking because point frames can represent a broader range of information (e.g., nonintentional points) than narrative frames.

Bibliography

Bareiss, R., & Osgood, R. (1993). Applying AI models to the design of exploratory hypermedia systems. In The Proceedings of the Fifth ACM Conference on Hypertext, (pp. 94-105). Seattle, WA: ACM Press.

Cleary, C., & Bareiss, R. (1994). A descriptive model of question asking during story acquisition interviews. In Proceedings of the Sixteenth Annual Conference of the Cognitive Science Society, . Atlanta, GA:

Cleary, C., & Bareiss, R. ([in preparation]). Automating Linking in Structured Hypermedia Systems. .

Collins, A., Brown, J. S., & Newman, S. E. (1983). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Eds.), Knowing, learning, and instruction: Essays in honor of Robert Glaser (pp. 453-494). Hillsdale, NJ: Lawrence Erlbaum Associates.

Osgood, R. (1994) The conceptual indexing of conversational hypertext. Ph.D., Northwestern University.

Osgood, R., & Bareiss, R. (1993). Automating index generation for constructing large-scale conversational hypermedia systems. In Proceedings of the Eleventh National Conference on Artificial Intelligence, (pp. 309-314). Menlo Park: AAAI Press.

Schank, R. C. (1977). Rules and topics in conversation. Cognitive Science, I, 421-441.

Schank, R. C., & Abelson, R. P. (1977). Scripts, Plans, Goals, and Understanding. Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R. C., & Cleary, C. (1995). Engines for Education (CD-ROM). In Hillsdale, NJ: Lawrence Erlbaum Associates.

Schank, R. C., Collins, G. C., Davis, E., Johnson, P. N., Lytinen, S., & Reiser, B. J. (1982). What's the point. Cognitive Science, I-6, 255-275.

Acknowledgments

We thank Chris Wisdo for his helpful comments on this paper and Carolyn Majkowski for assisting with the analysis of narrative linking. This work was supported by in part by ARPA (grant number # N00014-93-1-1212). The Institute for the Learning Sciences was established in 1989 with the support of Andersen Consulting. The Institute receives additional support from Ameritech and North West Water, Institute Partners, and from IBM.