

Event Concept Coherence in Narrative Text

Richard Alterman
Department of Computer Science
University of Texas at Austin
Austin, Tx 78712

What It Is.

For the purposes of this paper we will take the meaning of a concept to be its position in a structured network of mutually defining discrete symbols. For example, the concept 'give' is defined by its relationships to concepts like 'have' (i.e. to 'give' one must first 'have' the thing given, as a result of 'giving' there is a new 'having'), 'exchange' (i.e. 'giving' is a kind of 'exchange'), 'trading' (i.e. 'giving' and 'trading' are both kinds of 'exchange'), etc.

'Event/state concepts' are taken to be the words (terms) in a language which designate events or states. Two event/state terms, as they appear in text, are 'concept coherent' if their corresponding positions in a structured net of meaning are *proximal*. That is, given the text, "John gave Mary the bicycle. She has it.", the two event/state descriptions used in the text are 'concept coherent' because of the close inter-relationship between the concepts 'giving' and 'having'.

The point of an 'event concept coherence' representation of text has to do with the difficulty of establishing meaning representations. The idea is to capture only some of the properties of the text's meaning, to organize the text for further interpretation without losing any of the information of its original form. 'Event concept coherence' allows the text to be structured not so much on the basis of its meaning, but on the basis of a *property* of its meaning, its concept coherence. Because the proximity of concepts (i.e. concept coherence) is in some ways independent of the complexity of the relationships among terms in the network, it is possible to relax, at first, the content of the network and represent the connectivity of event/state concepts in terms of relationships easier to explain and establish. I will argue that initial structure for event concepts should be temporal relations, saving the causal, affect, and speech act interpretations for later stages.

The analysis of text in terms of event/state concept coherence results in a grouping of text into a hierarchy of chunks, each chunk centered around some event/state concept. Consider the following text adapted from "The Tale of the Pig" (Protter, 1961):

The pig trotted towards the stream carrying a bundle of clothes. The animal expertly soaked and scoured the laundry. The pig hung the clothes in the sun to dry. The pig gathered her laundry and trotted home.

The first set of terms which seem to cohere are 'trotting' and 'carrying'; 'carrying' involves 'travelling', which involves 'moving', and a kind of 'moving' is 'trotting'. The 'gathering' and 'trotting' described in the last sentence also seem to group around the concept 'carrying'; to 'carry' one must 'have' that which is carried, and 'gathering' results in a 'having'. 'Hanging clothes in the sun' is a way of 'drying', which, along with 'soaking', and 'scouring' are all parts of 'cleaning'. Finally, the 'cleaning' intersects the two 'carrying's' (i.e. the pig is 'moving' the laundry to the stream to 'clean' it and home again after it is 'cleaned'), so the text can be collected into a single chunk centered around the concept 'cleaning' (see figure 1).

A Little Structure

The ability to establish 'concept coherence' between event descriptions in text is dependent on the availability of a 'dictionary' (i.e. a structured network of concepts). In fact, an explanation (representation) of the concept coherence of the event descriptions in a piece of text can be derived by copying the relevant portion of the dictionary's structure. Because the dictionary is used for an initial structuring of the text, ideally its structure is simple, general, and useful.

NEXUS, the system developed to test this theory, uses a dictionary of between 100 and 150 event/state concepts interrelated by seven concept coherence relations (see figure 2). There are three taxonomic relations; one is a property inheritance relation, and the other two whole/part relations. The taxonomic relations are used to organize individual concepts into a hierarchical structure. Four of the relations are temporal; these are used to chain together sequences of typically co-occurring events and states. *Class/subclass* (sc) is the property inheritance coherence relation. In text, to find a connection between a pair of event/state terms it is frequently necessary to infer the inherited properties of either or




Figure 1: The pig cleans the laundry.

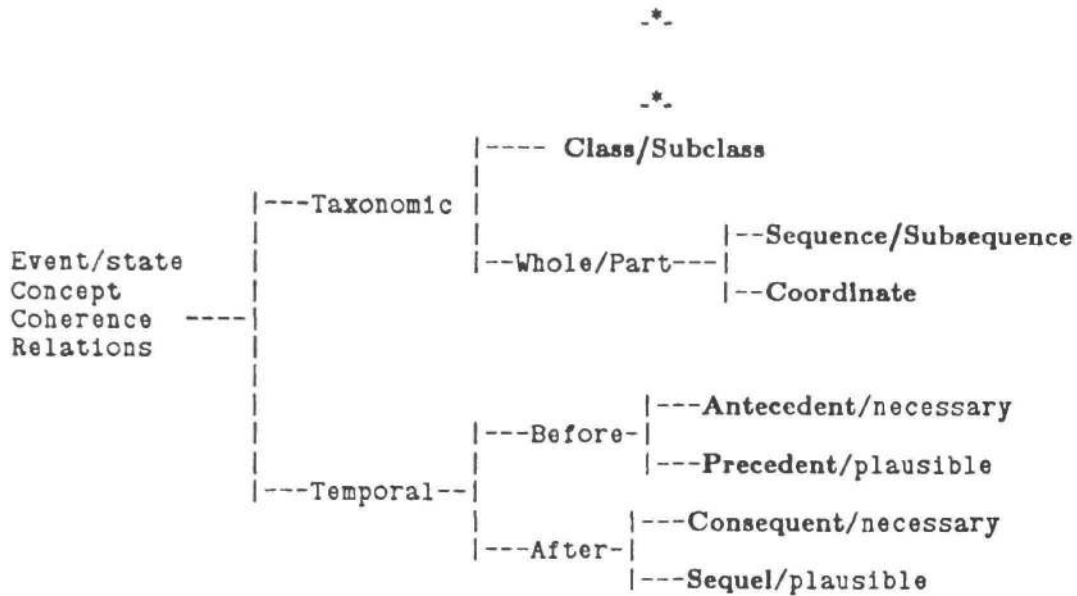


Figure 2: Characterizing the Relations

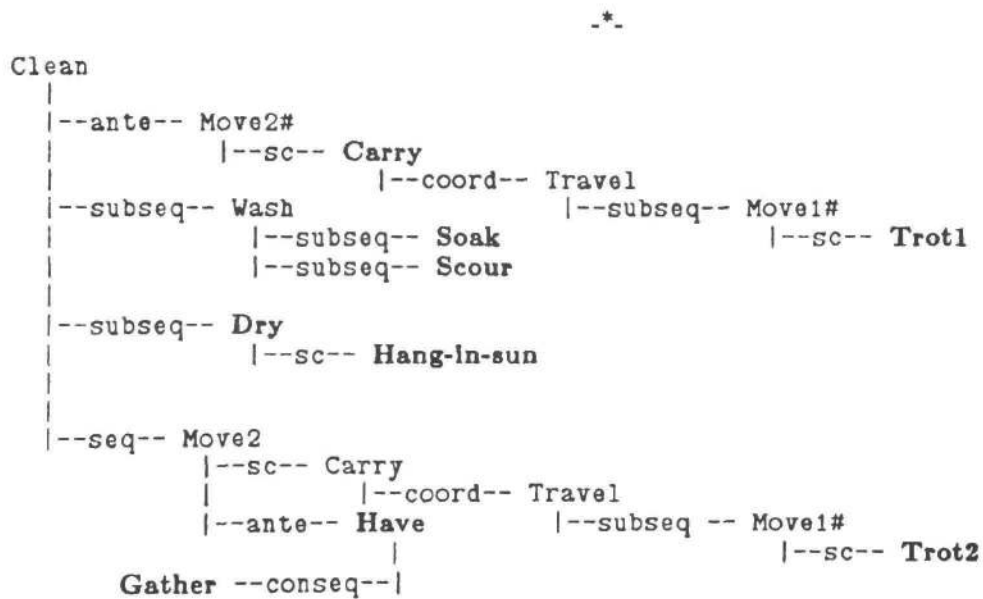


Figure 3: The Tale of the Pig Revisited

-*

both terms. Consider the pair of sentences: "John has the tickets. He snatched them from Bill." There exists a direct relationship between 'having' and 'taking' (i.e. as a result of 'taking something' one 'has it'). Because of the class/subclass relationship between 'taking' and 'snatching' NEXUS is able to access this relationship and represent the coherence of the sentences.

Experiments with NEXUS showed that class/subclass was the most frequently occurring coherence relation. This is not surprising considering the economy of using subclass terms in text; a subclass term simultaneously refines the meaning of its parent class while maintaining a implicit access to all its coherence relationships.

The *sequence/subsequence* (subseq) and *coordinate* (coord) relations can be differentiated by their temporal intervals. If one event is a part of another event, and it occurs for a subinterval of time, then the corresponding concepts are in a *sequence/subsequence* relationship. Consider the event descriptions: "John farms an acre of land. He'll be planting in the spring." A sub-activity of 'farming' is 'planting', another is 'harvesting'. Whenever the concept 'farming' applies the concept 'planting' applies for a subinterval of the 'farming' event, consequently 'farming' and 'planting' have a *sequence/subsequence* relationship.

If an event has parts that co-occur over the same time interval, then the corresponding event concepts are in a *coordinate* relationship. For example, "John carried the book. He walked, holding it in his hands.", every event which conveys a sense of 'carrying' in part also conveys a sense of 'moving' and 'holding' throughout the duration of the carrying, and there are 'moving' and 'holding' events that are not part of a 'carrying', and no 'carrying' event can completely be described by 'moving' or 'holding' - thus 'moving' and 'holding' are coordinates of 'carrying'.

Antecedent (ante), *precedent* (prec), *consequent* (conseq), and *sequel* (seq) are all temporal coherence relations. There are two ways of splitting the temporal relations. Antecedent and precedent concepts come before an event, and consequent and sequel concepts come after an event. The temporal relations can also be divided into logical and plausible groups of relations. Antecedent and consequent related concepts are logically related. Precedent and sequel concepts are plausibly connected.

If one event must necessarily occur before another event, the relationship between their corresponding concepts is classified as *antecedent*. For example, "John had some food. He ate it.": to eat one must first 'have' the thing which is to be eaten. Thus an antecedent of 'eating' is 'having food'.

If one event always, necessarily, occurs immediately after the other, then the relationship between their corresponding concepts should be marked as *consequent*. Take the following example, "John gave Mary a red kite. Mary has a red kite." Events of 'giving' are necessarily immediately followed by states of 'having'.

If one event, with some regularity occurs before another event, the relationship between their corresponding concepts can be classified as *precedent*. For example, "John opened the door. He entered.", events which can be described as 'entering' are sometimes preceded by sequences which enable the 'entering' and can be described as 'opening'.

If one event follows another with some regularity, the relationship between their corresponding concepts is *sequel*. An example of a sequel relationship is: "John cleaned the laundry. He carried it home." Sequences of 'cleaning' are typically followed by sequences where the laundry is 'moved'.

Although causal relations would have better described the relationships between two concepts, NEXUS's dictionary uses temporal relations because they are easier to establish. Recall that earlier I said that by relaxing the content of the net some of the complexity of establishing an initial interpretation of the text would be mitigated. Consider the use of the concepts 'cleaning' and 'carrying': "The pig cleaned the laundry. She carried it home." The causal relationship between these two event descriptions is not at all clear; it is neither result, nor enablement or reason. Using temporal relations it is easier to describe the relationship; it is not necessary that after the pig cleans the laundry she carried it home, but it is highly likely (i.e. sequel). A similar case can be made for the text: "Wild Bill rode to town. He tied his horse to the hitching post in front of the saloon." The causal relationship between 'riding' and 'tying a horse to a hitching post' is not simple, but with temporal relations it is easy to call the relationship sequel.

Furthermore, because temporal relations are causally neutral some of the problems associated with context can be finessed. Consider an example discussed by Wilks (Wilks, 1977) (p244): "The rock fell off the cliff and crushed John's lunch. Peter pushed it." The chain from 'pushing' to 'crushing' can be interpreted in one of two ways as either causal or goal. If it is later learned that Peter was angry at John then the chain is a goal chain. But if, instead, it is learned that John was clearing a spot to pitch his tent, the chain is a causal one. A system which produces causal chains would commit to one interpretation or the other (in Wilks' case the goal one), and potentially have to backtrack. NEXUS, because its dictionary uses temporal relations can produce a representation. At a later stage, NEXUS'

initial structuring of the text can be used as a basis for constructing a causal interpretation.

The Tale of the Pig Revisited.

Let us return to the example from "The Tale of the Pig", in this case adding some structures which a concept coherence representation would inherit from NEXUS' dictionary (see figure 3).

Consider the first 'carrying' event. Coordinates of 'carrying' are 'holding' and 'travelling'. Subsequences of 'travelling' are 'departing', 'moving', and 'arriving'. A subclass of 'moving' is 'trotting'. So NEXUS finds in its dictionary the path from 'carry' to 'trot' via the concepts 'travel' and 'move'.

Similarly the second 'carrying' event can be constructed. In this case, to connect 'carrying' to 'gathering', the 'carrying' must inherit from a second sense of 'moving' (i.e. John moved the box) an antecedent relationship to 'have' which is a consequent of 'gathering'.

The subsequence relations can be used to collect together the parts of 'cleaning'; subsequences of 'cleaning' are 'washing' and 'drying'. The subsequence relation is also used to collect together two parts of 'washing'; 'soaking' and 'scouring' are two subsequences of 'washing'. The 'hanging in the sun' is connected to 'drying' via a subclass arc. Finally the first 'carrying' is connected to the 'cleaning' via a 'moving#2' by an antecedent relation, and the second 'moving#2' by a sequel arc.

There are a few noteworthy features of this representation. In itself the representation shows the events of the story are coherent; the events can be collected together under the concept 'cleaning'. So the representation is de facto evidence of the text's coherency.

Notice that the analysis has produced two interconnected trees. In general an analysis of a piece of text will produce several interconnecting trees (one per major concept). The top node of a tree represents a summary of a concept. In summarizing this text, NEXUS deletes the tree centered around 'gathering', because it is 'preparatory', and produces the summary "The pig cleaned the laundry at the stream".

In the course of deriving the representation NEXUS resolves several references. For example it determines that 'animal' refers to 'pig', and that 'laundry' refers to 'clothes'. It also infers elided information; e.g. it infers that the pig cleaned the laundry at the stream.

NEXUS uses the representation as a basis for suggesting answers to several types of questions. The subsequence relationships between 'washing' and

'cleaning' and 'drying' and 'cleaning' is used to answer the questions, "Why did the pig wash the laundry?" and "Why did the pig dry the laundry?", with the statement, "She was cleaning the laundry." The subsequence relationships between 'washing' and 'cleaning' and 'drying' and 'cleaning' are also used by NEXUS to answer an 'how' question; "How did the pig clean the laundry? By washing and drying it." The combination of consequent and precedent relationships, subject to further semantic checks, can be used to answer 'goal' questions. For the question "Why did the pig gather the laundry?", NEXUS uses the coherence relations to suggest the answer "So she could carry it home."

Some Results and Conclusions

Event concept coherence should not be confused with the *discourse coherence* described by Hobbs (Hobbs, 1978). 'Event concept coherence' is primarily concerned with the coherence aspect of event/state concept definitions, and secondarily with how it is reflected in text. 'Discourse coherence' characterizes the text as discourse; it models the "coherent continuation moves" of the author or speaker (Hobbs, 1979) p68). 'Event concept coherence' represents the text by the dictionary relationships among the event concepts used in the text. With the exception of two *occasion* relations, enable and cause, 'discourse coherence' uses rhetorical relations, like example, parallel, elaboration and contrast. 'Event concept coherence' includes in its representation an explanation of the connection between two concepts, i.e. the inference path it found in the dictionary, 'Discourse coherence' does not. 'Event concept coherence' is derived in a data-driven fashion; 'discourse coherence' works from the top down looking for one of a set of rhetorical relations. 'Event concept coherence' has a dictionary of event concepts as its major source of knowledge; 'discourse coherence' requires a synthesis of semantic, factual, expert, stylistic, and grammatical knowledge. These comments are meant to clarify the many differences that exist between 'event concept coherence' and 'discourse coherence', not to detract from Hobbs' work.

The event concept coherence representation system is realized in a program called NEXUS. The subsystem which produces the representations is called TRACE, the question answerer QUEST, and the summarizer SUM. The system is programmed in procedural logic using HCPVR (Chester, 1980a, Chester, 1980b).

Concept coherence analysis has been applied to text previously accounted for by scripts (Schank & Abelson, 1977, Cullingford, 1978), plans (Schank & Abelson, 1977, Wilensky, 1978, Wilensky, 1981), story trees (Rumelhart, 1975), schema-narrative trees (Simmons, 1982, Simmons, 1983, Correia, 1980),

and speech act (Allen, 1979, Cohen and Perrault, 1979, Cohen, Perrault, and Allen, 1981) theories. Event concept coherence is not intended to replace these top-down theories, but more to complement them with a theory that attempts to grow the more complex interpretations of text from the bottom-up.

TRACE has successfully been applied to eight samples of text. It is important to remember that for each of the eight examples TRACE used the identical dictionary of concepts. Three of the samples came from the AI literature; these included stories which had been accounted for by either scripts (Lehnert, 1977), plans (Wilensky, 1978) or story trees (Rumelhart, 1975). The other five samples, including a more difficult version of "The Tale of the Pig" example, come from a book of folktales (Protter, 1961). For further details see Alterman (Alterman, 1982).

Acknowledgements

I would like to thank Bob Simmons and Elaine Rich for their comments on an earlier draft of this paper.

References

- Allen, James. *A plan-based approach to speech act recognition*. PhD thesis, University of Toronto, 1979.
- Alterman, Richard. *A system of seven coherence relations for hierarchically organizing event concepts in text*. Technical Report TR-188, University of Texas at Austin, 1982.
- Chester, Dan. *Using HCPRVR*. Technical Report, The University of Texas at Austin, 1980.
- Chester, D. *HCPRVR: a logic program interpreter in LISP*. AAI, 1980.
- Cohen, Phillip R. and Perrault, C. Raymond. Elements of a plan-based theory of speech acts. *Cognitive Science*, 1979, 3(4), 177-202.
- Cohen, P. R., Perrault, C. R., and Allen, J. F. *Beyond Question-Answering*. Technical Report, Bolt Beranek and Newman Inc., 1981.
- Correia, Alfred. Computing story trees. *American Journal of Computational Linguistics*, 1980, 6, 135-149.
- Cullingford. *Script application: computer understanding of newspaper stories*. PhD thesis, Yale University, 1978.
- Hobbs, Jerry. *Why is discourse coherent?* Technical Report Technical Note 176, SRI, 1978.
- Hobbs, Jerry. Coherence and coreference. *Cognitive Science*, 1979, 3(1), 67-90.
- Lehnert, Wendy. *The Process of Question Answering*. PhD thesis, Yale University, 1977.
- Protter, Eric (Ed.). *A children's treasury of folk and fairy tales*. : Channel Press Inc 1961.
- Rumelhart, David E. Notes on a schema for stories. In Bobrow, Danial G. and Collins, Allan (Eds.), *Representation and Understanding*, : Academic Press, Inc, 1975.
- Schank, R. C. and Abelson, R. P. *Scripts, plans, goals, and understanding*. : Lawrence Erlbaum Associates 1977.
- Simmons, Robert F. A narrative schema in procedural logic. In Clark, K. L. and Tarnlund, S.-A. (Ed.), *Logic programming*, : Academic Press, 1982.
- Simmons, Robert F. *Computations from the english*. : Prentice-Hall 1983. In Press.
- Wilensky, R. *Understanding goal-based stories*. PhD thesis, Yale University, 1978.
- Wilensky, R. PAM. In (Ed.), Schank, R. C. and Riesbeck, Christopher (Eds.), *Inside computer understanding: five programs plus miniatures*, : Lawrence Erlbaum Associates, 1981.
- Wilks, Yorick. What sort of a taxonomy of causation to we need? *Cognitive Science*, 1977, 1, 235-264.

