Affect and Memory Representation

Wendy G. Lehnert

Department of Computer Science
Yale University

1. Affective Knowledge Structures and Computers

Many philosophers, psychologists, and A.I. workers have taken various positions on the issue of machines and emotion. Some argue that a computer can never "experience" human emotions in any significant sense because it just doesn't make sense to attribute conscicusness to an inorganic and programmed system [Gunderson 1971, Puccetti 1968, Scriven 1960, Ziff 1959] Others argue that our subjective sense of emotional experience is too "intuitive" and ill-defined a candidate for computational modelling [Dreyfus 1972, Weizenbaum 1976]. Still others argue that emotion will be a natural and necessary consequence of intelligent information processing, an inevitable side-effect of intelligence [Kenny 1963, Simon 1967, Doyle 1980, Sloman 1981]. And then there are always "rational purists" who consider emotional experience totally irrelevant to reasoning processes and therefore of no consequence to artificial intelligence whatsoever.

Many lines of reasoning have been invoked to secure these various positions [Dennett 1978], although most of the arguments are conducted with a distinctively philosophical tone. It is ironic that the most passionate advocates in these debates rarely argue from first-hand experience with computer simulations. Why does A.I. seem to be so silent on the the topic of emotion and computers? I cannot speak for everyone in the field, but I would guess that a lot of us prefer to avoid the whole morass because we believe that the questions being answered are not the questions we should be asking.

A computer can have knowledge of human emotionality in the same sense that it can have knowledge of mass spectroscopy, medical diagnostic techniques, or payroll data. Computers do not have to "be" emoting entities to use this knowledge any more than they have to "be" chemists, physicians, or bureaucrats to use knowledge specific to those professions. If it is difficult to give knowledge of emotions to computers, it is only difficult for the same reason that a thousand other topics are difficult for computers: people do not have a rigorous understanding of their intuitive knowledge in terms of information processing requirements. We need to develop (1) symbolic systems of internal representation, and (2) processing strategies to manipulate these symbolic structures. These two requirements are universal to all A.I. efforts, and the difficulties involved are not significantly amplified when the knowledge to be encoded is knowledge of human emotionality.

When we apply knowledge of emotions to an information processing task, we can evaluate our expertise on emotions by evaluating the overall effectiveness of the larger information processing system. What experience has A.I. had with affective knowledge structures? Our experience is admitedly limited, but it is not totally non-existent. To date, three distinct task orientations have touched on affective manipulations of one sort or another.

1) belief system maintainance
2) conversational simulations
3) narrative text processing

An ambitious implementation of belief system maintainance was attempted by Kenneth Colby in the early 60's [Colby 1967; 1973]. While Colby is best known for PARRY, the paranoid conversationalist [Colby 1975], his earlier work was aimed at a more general simulation of neurotic thought processes [Colby 1963; 1965]. Colby was specifically interested in simulations of Freudian defense mechanisms when they surface in clinical dialogues between psychiatrists and patients. His work involved affective manipulations, but only in a very superficial sense. Colby utilized "emotion monitors" which were numerical parameters with names like "excitation," "self-esteem," "danger," "well-being," and "pleasure." While Colby's simulations were never intended to implement a complete system of affective representation, he nevertheless found it necessary to maintain and manipulate these numeric parameters. For example, the "excitation" monitor reflected the overall anxiety of the system - a factor that any psychotherapist would want to take into account. Whether or not someone's anxiety level can be adequately represented on a scale of 1 to 10 is another question.

It is inevitable that belief system manipulations manifest themselves most naturally in interactive conversation. Colby was forced into conversational task orientations when he began his work on belief systems, and this eventually drew him toward PARRY. PARRY also utilized numeric parameters for "anger," "fear," and "mistrust," - a somewhat more narrow set than was needed for generally neurotic simulations. While PARRY is the only conversational system that I know of which has implemented an affective component, it is clear that any conversational system would require affective manipulations if it was designed to simulate emotional responses [Schank and Lehnert 1979].

Although Colby was primarily interested in thought processes, his simulations became thoroughly mired in language processing difficulties. Colby's sentence processing techniques relied on lexical pattern matching routines, and his internal memory representations were lexically-oriented as well. These devices were ineffective substitutes for natural language processing strategies, and Colby's models were significantly hampered by inadequate representational techniques [Boden 1977]. Similar impediments were encountered by other researchers who tackled belief systems early on [see e.g. Abelson 1973], so it comes as no surprise to see that the most recent work on belief systems is thoroughly grounded in theories of natural language processing and internal memory representation. [Carbonell 1978]. Whatever one's ultimate research goal (models of belief systems, memory organization, etc.) all dialog simulations are primarily natural language processing systems, and any attempt to circumvent this fact is destined to fail. In fact, when the research goal is a model of human emotion,

This work was supported in part by the Advanced Research Projects Agency under contract N00014-75-C-1111 and in part by the National Science Foundation under contract IST7918463.

it is far better to embrace the challanges of natural language processing with open arms: a natural language processing project provides the only naturalistic and realistic laboratory for theories of affective memory representation.

For example, narrative texts provide a rich proving ground for theories of affect - one need only look at stories that involve emotional reactions and emotional behavior. Yet affective knowledge representation has not been systematically tackled within natural language processing programs until very recently. The remainder of this paper will outline some of my own experiences in trying to implement affective knowledge structures in a narrative text processing system, BORIS.

The BORIS system currently utilizes a limited system of affective representation and affect-related inference mechanisms. In addition to these representational techniques for affects per se, the recently proposed TAU knowledge structure (Thematic Affect Unit) contributes to BORIS's affective inference capabilities as well [Dyer 1981]. Since descriptions of BORIS's processing techniques can be found elsewhere [Dyer and Lehnert 1980, Lehnert, Dyer, Johnson, Yang and Harley 1981], we will not go into a description of BORIS here - we will instead take this opportunity to explain how a computational model of affect might contribute to other models of affect that are not computational in nature.

To begin with, a computational model must address a specific set of problems to be solved. In language processing, affective manipulations are needed for four general inference situations:

(1) Given an event description, infer an affective antecedent:

"John took a valium." (=> John is upset)

(2) Given an event description, infer an affective consequent:

"John got a big raise." (=> John is happy)

(3) Given an affective state, infer its likely antecedent:

"After the hurricane, John was depressed."
(John suffered a loss => John is depressed)

(4) Given an affective state, infer its likely consequence:

John was so happy about his royalty check, he made reservations at Reno Sweeny's.
(... => John intends to celebrate)

In many cases, we must combine two or more of the above inference types to make sense of an implicit causality:

"After the hurricane, John saw a shrink."

To see how this process-orientation differs from a purely psychological approach to the problem, we will look at some inferences problems in a narrative text, and see how far a non-process-oriented model can go in helping a system like BORIS. When we first began to look at affect in BORIS, we were greatly inspired by Ira Roseman's model for representing affective states [Roseman 1979]. There are of course other approaches to affect [deRiveria 1977, Izard 1977], but we will not attempt to survey all the relevant proposals here. Readers who are familiar with alternative systems can judge for themselves whether similar troublespots would arise in trying to implement another system.

2. Conceptual Decomposition for Affective States

In the Roseman system, emotional states are represented by decomposition into five dimensions. Four of the dimensions assume positive and negative fields, while the fifth assumes a three-valued spectrum:

Five Dimensions of Affect

1) Motivational Status (desirability) [+,-]
2) Situational Status (attainment) [+,-]
3) Probability Status (certainty) [+,-]
4) Legitimacy Status (deservedness) [+,-]
5) Agency Status [self, other, circumstantial]

When an event is mapped into its appropriate place on each of the five spectrums, we can predict emotional reactions to the event. For example, (a) wanting a ticket to a sold-out Grateful Dead Concert describes a mental state with a positive motivation (wanting it) and a negative situation (not having it); (b) winning a ticket to a Grateful Dead Concert is an event with a positive motivation (wanting it) and a positive situation (having it); (c) losing the ticket has negative motivation (not wanting to lose it) and positive situation (having lost it); and (d) finding it again involves a negative motivation (not wanting it lost) with a negative situation (not having it lost). If all of this happens circumstantially, we expect to see (a) sorrow, (b) joy, (c) distress, and (d) relief. Using all five dimensions, Roseman's system differentiates 13 primary emotions. These are listed below with a vector encoding of the five dimensions as listed above. For example, (+ + + - S) corresponds to a positive motivation, positive situation, positive probability, negative legitimacy, and self-agency. An "*" indicates that the corresponding dimension can assume any value.

PRIMARY EMOTIONS

M S P L A		M S P L A	
(+ + + * C)	JOY	(+ + * * O)	LIKING
(+ + - * C)	HOPE	(- - * * O)	LIKING
(- - - * C)	HOPE	(+ - * - O)	DISLIKING
(- - + * C)	RELIEF	(- + * - O)	DISLIKING
(- + + - C)	DISTRESS	(- + * + O)	ANGER
(- + * + C)	FRUSTRATION	(+ * + O)	ANGER
(+ - * + C)	FRUSTRATION	(+ - - - C)	FEAR
(+ - + - C)	SORROW	(- + - C)	FEAR
(+ + * + S)	PRIDE	(+ - * + S)	REGRET
(- - * + S)	PRIDE	(- + * + S)	REGRET
(* * * - S)	GUILT		

Many lexical descriptions of emotionality are used to reference more than one conceptual configuration. For example, John could "regret" flunking a test (- + + + S). Alternatively, if John got a high B on the test, he might "regret" not getting an A (+ - + + S). These are two distinct senses of regret: we can regret what happened, and we can regret what didn't happen. People who dwell on (- + + + S) configurations kick themselves for past mistakes while people who dwell on situations involving (+ - + + S) are melancholy dreamers. We can describe either personality in terms of past

regrets, but important conceptual distinctions are lurking beneath these words.

Lexical ambiguities at the conceptual level make it difficult to describe Roseman's 13 primary emotions to everyone's satisfaction. For example, one could argue that "liking" is not an emotion at all but an attitude. The appropriate emotion for (+ + * * 0) and (- - * * 0) is really one of gratitude. Or perhaps "distress" should be called "discomfort." It is instructive to engage in this sort of criticism as an intuitive exercise, but a better way to test Roseman's system is with a computer implementation.

3. Implementing Affect

When we tried to implement Roseman's system in BORIS, we ran into some interesting difficulties. To get a sense of these, we will look at a sample text that BORIS processes, highlighting some problem areas encountered.

****

A BORIS Narrative

Richard hadn't heard from his old roommate Paul for years. Paul had loaned Richard money which was never paid back, but now he had no idea where to find his old friend. When a letter finally arrived from San Francisco, Richard was anxious to find out how Paul was.

Unfortunately, the news was not good. Paul's wife Sarah wanted a divorce. She also wanted the car, the house, the children, and alimony. Paul wanted the divorce, but he didn't want to see Sarah take everything he had. His salary from the state school system was very small. Not knowing who to turn to, he was hoping for a favor from the only lawyer he knew. Paul gave his home phone number in case Richard felt he could help.

Richard eagerly picked up the phone and dialed. After a brief conversation, Paul agreed to have lunch with him the next day. He sounded extremely relieved and grateful.

The next day, as Richard was driving to the restaurant, he barely avoided hitting an old man on the street. He felt extremely upset by the incident, and had three drinks at the restaurant. When Paul arrived, Richard was fairly drunk. After the food came, Richard spilled a cup of coffee on Paul. Paul seemed very annoyed by this so Richard offered to drive him home for a change of clothes.

When Paul walked into the bedroom and found Sarah with another man, he nearly had a heart attack. Then he realized what a blessing it was. With Richard there as a witness, Sarah's divorce case was shot. Richard congratulated Paul and suggested that they celebrate at dinner. Paul was eager to comply.

****

There are a number of important affect-related inferences in this story. For example, we should infer that Richard felt bad about spilling his coffee on Paul, and his offer to drive Paul home was motivated (at least in part) by a desire to alleviate guilt. In the next sentence, when Paul finds Sarah, we should not assume that Paul suffered a cardiac arrest; he is just very surprised. We must also understand why it made sense for Richard to congratulate Paul and suggest a celebration. What did Paul have to celebrate? Adulterous mates are not normally greeted with such enthusiasm, so the celebration must be causally connected to something else that Paul should feel good about. Notice that if Richard had expressed his heartfelt condolences to Paul instead of congratulating him, this would also make sense. Paul's affective state is complex and must be fully understood to accommodate these various possibilities.

To make affective inferences, BORIS needed to interpret events and states from the story in terms of Roseman's five affect dimensions. "Motivation" and "situation" were relatively easy to recognize by relying primarily on goal states. But the three remaining dimensions proved to be trickier than expected. We will look at some difficulties in "agency" recognition, although similar illustrations could have been chosen from "certainty" and "legitimacy" as well.

Initially, we thought that the agent for an event would simply be the physical actor of the event [Schank 1975]. We quickly discovered otherwise. For example, in a question answering task, experimental subjects indicated that Richard was happy to receive the letter from Paul. He wasn't grateful, and he didn't like Paul any more than before; he was simply happy. In order to infer that Richard was happy to get Paul's letter, we have to ascribe a circumstantial agent when Richard gets the letter. If the letter's arrival was encoded with other-agency, then Richard would either like Paul or feel grateful to Paul for getting the letter. Simple joy can only come from circumstantial agency. But the letter's arrival is encoded as an MTRANS event with actor = Paul (Paul sent the letter). If agency is not a function of an event's actor, what is it? BORIS was (and still is) stymied by the agency problem.

It seems that agency is a function of actors but more specifically, intentional actors. Notice how the affective inference changes if Richard believes that Paul sent the letter just because Paul wanted to make Richard happy. Now it is much more reasonable for Richard to like Paul or feel grateful to Paul for sending the letter. If X knowingly does Z to make Y happy, and Z succeeds in making Y happy, then Y will like X for doing it. If Paul does something intending to make Richard happy, then Richard experiences the event with other-agency. But if Paul does something which only makes Richard happy incidentally, then Richard experiences the event with circumstantial-agency. Knowledge of an actor's ultimate intentions is needed to establish affective agency for inference purposes.

In addition to intentionality, affective agency can be influenced by an actor's degree of social responsibility. For example, it makes sense that Paul got annoyed when Richard spilled coffee on him. But what is annoyance? Annoyance can be a variant of anger or dislike (Paul was annoyed with Richard), both of which require other-agency. Richard didn't intend to spill the coffee, but he was nevertheless responsible for the event (albeit innocently), and this responsibility gives us other-agency. Annoyance is even more ambiguous in the sense that it may also describe frustration, which involves circumstantial-agency: it is rotten luck to have someone spill coffee on you. If Paul is upset, but not upset with Richard specifically, then his annoyance is one of pure frustration.

Since "annoyed" is ambiguous, and this particular example could go either way, it is useful to look for limiting cases which force one interpretation over another. For example, it seems reasonable that Paul might be more annoyed with Richard for the accident since Richard was drunk. If Richard were sober, he would somehow be less at fault. Suppose a frail little old lady is carrying a cup of coffee, and as she passes by Paul she collapses from a heart attack. Do we expect Paul to be angry at the old lady for spilling coffee on him? Not likely. Now suppose a boisterous drunk lurches past Paul and drops a drink on him. Do we expect Paul to be angry at the drunk? Sure. Neither event was intended, but a drunk is more responsible for his actions than a heart attack victim. People choose to get drunk, but they don't choose to have heart attacks. The element of free will operating in a drunk renders him more responsible for his accidents: a drunk chooses to be accident-prone. Since BORIS has no heuristics for assessing relative degrees of responsibility, BORIS defaults to circumstantial agency and therefore interprets Paul's annoyance as one of pure frustration. This is not altogether right, but a more correct interpretation requires an assessment mechanism for social responsibility.

One final problem with agency involves events that cause complex affect states. When Paul catches Sarah in their bedroom with another man, he witnesses and reacts to an event involving two actors. The event is assumed to be intentional (we are given no reason to interpret the bedroom activities as a rape) and at least one of the lovers must be responsible for it. So it seems that we have a clear-cut candidate for other-agency. Since this event will save Paul from a nasty court battle and divorce settlement, it is desirable from Paul's perspective. Sarah's activity can therefore interpreted by Paul as a desirable, positively attained, certain, and illegitimate (she's violating their marital contract) event of other-agency (+ + + - 0). But this configuration brings us to the improbable prediction that Paul will like Sarah and her lover, or feel grateful to them for engaging in their illicit activity.

The difficulty with this example is the complexity of Paul's emotional state. He may be happy about the settlement implications, but he is probably very unhappy about his territorial rights. Even if he doesn't feel possessive about Sarah (Paul did say that he also wanted a divorce), he has a right to feel put out by a stranger in his bedroom, to say nothing of his bed. His privacy is surely being violated on at least one level, and we are assured of his negative reaction when we are told that he "almost had a heart attack." So Paul's reaction is mixed: it has a strong negative component (- + + - 0) and a more far-sighted positive component (+ + + - 0) as well. This explains why it would make sense for Richard to either express sympathy or offer congratulations. It seems most appropriate to first offer condolences and then congratuations, but either one can be understood as a reasonable reaction on Richard's part.

Special heuristics must be invoked for complex emotional states, and higher level knowledge structures will be needed to handle inferences in these cases. For example, the representational system of plot units (which grew out of our experience with BORIS's affect analysis), includes a special structure called "hidden blessing" to handle situations like Paul's reaction to Sarah [Lehnert 1980, 1981a, 1981b]. The hidden blessing plot unit encodes any event that causes an initial negative reaction which later yields to a dominately positive emotion. A similar "mixed blessing" plot unit handles cases where the initial reaction is positive, but a negative emotion follows.

In general, a plot unit is a configuration of three affect states: (1) "M" mental states with neutral affect, (2) "+" events that cause positive affects, and (3) "-" events that cause negative affects. Each affect state is interpreted with respect to a specific character, although plot units may contain multiple affect states that involve more than one character. For example, the retaliation unit involves two characters and five affect states:

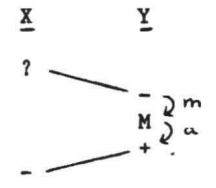

This configuration tells us that X did something (?) which caused a negative reaction in Y (-). This negative event motivated a mental state (M) in Y, which was subsequently actualized by a positive event (+) for Y, and a negative event (-) for X. In other words, X did something that distressed Y, so Y retaliated by doing something to distress X. The vertical and diagonal links in this diagram describe various causal relationships between affects states within characters and across characters.

Affect state maps are constructed for each character in a story, as a way of tracking that character's emotional ups (+'s) and downs (-'s), and specific plot units are recognized when the linkages across affect states indicate that a given plot unit configuration has been encountered. If X were to get back at Y for Y's retaliation, we would have two instances of the retaliation plot unit sharing two common affect states:

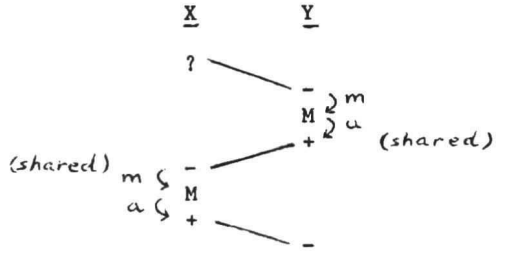

A plot unit graph for a story can be generated by creating a graph node for each instantiated plot unit, and placing arcs between all pairs of plot units that share at least one common affect state. The above affect state map yields a graph of two nodes connected by an arc, while the BORIS divorce story involves 24 plot units in a connected graph structure.

It appears that the connectivity features in a plot unit graph provide a strong basis for summarization algorithms. When a story's plot unit graph contains a pivotal node (a node with maximal connectivity), we can expect a short summary of the story to be based on the conceptual content of that pivotal node. For example in the BORIS divorce story, the hidden blessing unit is pivotal, and we can summarize the story by saying "Paul saved himself from a nasty divorce settlement when he

accidentally found his wife in bed with another man." The hidden blessing unit encodes Paul's discovery as an event of mixed affect states, and a synopsis of any hidden blessing has to explain what was ultimately good about an initially negative event.

We have been experimenting with plot units primarily as a basis for narrative text summarization, [Lehnert 1980, 1981a, 1981b; Lehnert, Black, and Reiser 1981; Reiser, Lehnert, and Black 1981], but their use as a predictive knowledge structure for affective inference remains to be explored. Initial efforts in this direction led to the development of a slightly higher level of memory representation (Thematic Affect Units) that relates to adages and fables [Dyer 1981].

4. Conclusions

Our experience with Roseman's affect analysis and BORIS suggests that affective inferences are dependent on a substantial range of other inference mechanisms. It is not possible to study problems of affect without addressing seemingly unrelated problem areas. The current state of the art in language processing allows us to tackle recognition techniques involving scripts, goals, plans, interpersonal themes, plot units, and thematic affect units, all of which can contribute to affect recognition techniques. But affect analysis can also lead us into largely uncharted regions of intentionality and social responsibility, just to name two areas we've discussed.

We have not attempted to compile a list of all the related knowledge needed to handle affect, because this list is likely to be a comprehensive list of all knowledge structures used for language processing. Interestingly enough, there will probably be no knowledge structures devoted exclusively to affect. One could argue that the presence of a Roseman-like vector (+ - + + S) within computational memory constitutes an affect-specific knowledge structure. But this structure is just a processing artifact: we do not expect to find these vectors in long or short-term memory representations. If we needed to explicitly encode "John was angry," we might be reduced to vector notation, but as soon as this sentence is embedded in a context which tells us what happened to John and why he feels angry, his anger will likewise be embedded in some larger structure. TAU's fill this role already, and plot units (originally called "affect units") also operate along these lines. For example, vindictive or vengeful feelings can be readily derived from the retaliation plot unit.

In conclusion, I would say that it is altogether too early to pass judgement on any specific representational techniques for affects. Roseman has supplied the computational environment with a valuable framework in which to work (or perhaps play) and other psychological theories may also prove to be valuable, although I know of no others that have been (even partially) implemented. This particular area is a difficult one for A.I. implementations because it draws on so many other complex problem areas in memory and cognition.

It was nevertheless valuable to attempt an implementation of Roseman's model, if only for the sake of the spin-offs that emerged. By confronting problems of affect representation and affect-related inferences, we were led to the design of two new knowledge structures for narrative text analysis. Plot units and thematic affect units were natural devices for solving affect-oriented processing problems. The importance of affective knowledge structures for memory representation deserves further exploration in both the psychological and A.I. research paradigms. Recent research in both areas has uncovered some provocative results concerning emotion and memory which pave the way for further investigations. From the A.I. end of the world, we are seeing how the use of plot units for narrative text summarization suggests that emotional states of characters in a story play a far more central role in high-level memory representation than was previously suspected [Lehnert 1981a]. At the same time, psychologists are demonstrating how the moods and attitudes of experimental subjects can dramatically alter their patterns of memory retrieval [Bower 1981].

So affects appear to play many different roles in human cognition. In particular, we have seen how knowledge of emotional reactions can be crucial to various information processing tasks, including narrative text comprehension. Still a great deal of work remains to be done if we are going to fully understand the roles of emotional knowledge and experience in human information processing. We must turn to psychology labs and LISP programs for the answers to our questions, putting aside the more philosophical speculations about machines and emotions. It will be much easier to resolve our speculative debates in the face of some solid research, and the quality of our interdisciplinary dialogs will be greatly enhanced by empirical contributions from psychology and artificial intelligence.

BIBLIOGRAPHY

Abelson, R.P. (1973). "The Structure of Belief Systems." in Computer Models of Thought and Language. (eds: Schank and Colby). San Francisco: W.H. Freeman.

Boden, M. (1977). Artificial Intelligence and Natural Man. New York: Basic Books.

Bower, G. H. (1981). "Mood and Memory." American Psychologist. Vol. 36, No. 2. 129-148

Carbonell, J. G. (1978). "POLITICS: Automated Ideological Reasoning." Cognitive Science. Vol. 2, No. 1. 27-52.

Colby, K. M. (1963). "Computer Simulation of a Neurotic Process." in Computer Simulation of Personality: Frontier of Psychological Research (eds: Tomkins and Messick). New York: Wiley.

Colby, K. M. (1965). "Computer Simulation of Neurotic Processes." in Computers in Biomedical Research, Vol. 1 (eds: Stacy and Waxman). New York: Academic Press

Colby, K. M. (1967). "Computer Simulation and Change in Personal Belief Systems." Behavioral Science, 12, 248-253.

Colby, K. M. (1973). "Simulations of Belief Systems." in Computer Models of Thought and Language. (eds: Schank and Colby). San Francisco: W.H. Freeman.

Colby, K. M. (1975). Artificial Paranoia. New York: Pergamon.

Dennett, D. C. (1978). Brainstorms. Montgomery, Vermont: Bradford Books.

deRivera, J. (1977). A Structural Theory of the Emotions. New York: International Universities Press.

Doyle, J. (1980). "A Model for Deliberation, Action, and Introspection." AI-TR.581, MIT Artificial Intelligence Laboratory.

Dreyfus, H. L. (1972). What Computers Can't Do: A Critique of Artificial Reason. New York: Harper and Row.

Dyer, M. (1981). "The Role of TAU's in Narratives." Proceedings of the Third Annual Conference of the Cognitive Science Society. Berkeley, Calif.

Dyer, M. and Lehnert, W. (1980). "Organization and Search Processes for Narratives. Department of Computer Science Research Report #175. Yale University, New Haven, Conn.

Gunderson, K. (1971). Mentality and Machines. New York: Doubleday.

Izard, C. E. (1977). Human Emotions. New York: Plenum Press.

Kenny, A. (1963). Action, Emotion, and Will. London: Routledge and Kegan Paul.

Lehnert, W. (1980). "Narrative Text Summarization." Proceedings of the First Annual National Conference on Artificial Intelligence. Stanford, Calif.

Lehnert, W. (1981a). "Plot Units and Narrative Summarization." Cognitive Science. (in press).

Lehnert, W. (1981b). "Plot Units: A Narrative Summarization Strategy." in Strategies for Natural Language Processing. (eds: Lehnert and Ringle). New Jersey: Lawrence Erlbaum (in press).

Lehnert, W., Black, J., and Reiser, B. (1981). "Summarizing Narratives." Proceedings of the Seventh International Joint Conference on Artificial Intelligence. Vancouver, British Columbia.

Lehnert, W., Dyer, M., Johnson, P., Yang, C.J., and Harley, S. (1981). "BORIS -- An Experiment in In-Depth Understanding of Narratives." Department of Computer Science Research Report #188. Yale University, New Haven, Conn.

Puccetti, R. (1968). Persons: A Study of Possible Moral Agents in the Universe. London: Macmillan.

Reiser, B., Lehnert, W., and Black, J. (1981). "Recognizing Thematic Units in Narratives." Proceedings of the Third Annual Conference of the Cognitive Science Society. Berkeley, Calif.

Roseman, I. (1979). "Cognitive aspects of emotion and emotional behavior." presented at the meeting of the American Psychological Association, New York. (unpublished manuscript, Department of Psychology, Yale University, New Haven, Conn.)

Schank, R. (1975). Conceptual Information Processing. Amsterdam: North-Holland.

Schank, R. and Lehnert, W. (1979). "The Conceptual Content of Conversation." Proceedings of the Sixth International Joint Conference on Artificial Intelligence. Tokyo.

Scriven, M. (1960). "The compleat robot: a prolegomena to androidology." in Dimensions of Mind: A Symposium. (Ed: S. Hook). New York: New York Universities Press.

Simon, H. A. (1967). "Motivational and Emotional Controls of Cognition." (reprinted in Models of Thought. Yale University Press. 1979)

Sloman, A. (1981). "Why Robots Will Have Emotions." Proceedings of the Seventh International Joint Conference on Artificial Intelligence. Vancouver, British Columbia.

Weizenbaum, J. (1976). Computer Power and Human Reason. San Francisco: W.H. Freeman and Company.

Ziff, P. (1959). "The feelings of robots." Analysis. No. 19, 64-68.