Emotional influences in memory and thinking: Data and theory

Gordon Bower
Stanford University

In M.S. Clark & S.T. Fiske (Eds.), Affect and Cognition. Hillsdale, NJ: Lawrence Erlbaum Associates, 1982.

13 Emotional Influences in Memory and Thinking: Data and Theory

Gordon H. Bower
Paul R. Cohen
Stanford University

INTRODUCTION

We have investigated emotional influences on human memory, perception, judgment, and thinking. We find powerful effects of people's feelings upon their cognitive processes. Our first group of results show how a person's feelings act like a selective filter that is tuned to incoming material that supports or justifies those feelings; the filter admits material congruent with the perceiver's mood but casts aside incongruent material. Feelings cause congruent stimuli to become more *salient,* to stand out more, arouse more interest, cause deeper processing and greater learning of congruent material. This filtering is important insofar as it determines what gets stored in memory in the first place.

Second, we find that people's feelings affect what records they can retrieve from memory. People can best retrieve events originally learned in a particular mood by somehow reinstating or returning to that same mood. We have several demonstrations of this "mood-state-dependent retrieval."

Third, we find emotional influences upon thinking and judgments. People's social perceptions as well as their imaginative fantasies are subjective; they are easily influenced by their feelings at the moment. These influences occur when people evaluate their friends, themselves, their possessions, and their future.

We propose a preliminary theory to explain these empirical results. This theory assumes that memories of events are recorded into a semantic network, and that different emotions can be represented by different units or nodes in this same network. When active, an emotion will become associated with coincident events. Memories, concepts, and perceptual categories are retrieved by the spreading of activation from the current emotion unit as well as from the units corresponding to the explicitly presented retrieval cues.

Whereas this network theory suffices to account for our results in a general way, it is incomplete in not addressing how emotional reactions are evoked by events. We introduce the "blackboard" control structure to model how the person combines several knowledge sources in arriving at an emotional interpretation of a situation. This enables us to deal with certain theoretical puzzles such as reappraisals of earlier emotional experiences and hot versus cold uses of emotional terms.

Experimental Background. Most of our experiments use a common methodology that is quite effective. We induce emotional states (happiness, sadness, or anger) in our college-student volunteers via hypnotic suggestions. We first select volunteers who are highly hypnotizable. After hypnotizing them, we suggest that they place themselves in a specified emotional state by remembering an appropriate scene from their life and reliving that emotional experience in their imagination. Their happy scenes are typically moments of success or personal intimacy; sad scenes are usually moments of failure, loss, or rejection. After getting into the specified mood at a moderate level, subjects are asked to maintain it while they perform our experimental tasks. The tasks require 5 to 20 minutes, during which we typically remind subjects to refresh and maintain their mood. The hypnotic moods seem quite genuine and produce results comparable to mood effects obtained by naturalistic means. With this background, we now describe some of our results. Some of these were described earlier (Bower, 1981), although several new findings are reported here.

Selective Filtering and Learning

The basic hypothesis is that people's feelings cause certain environmental stimuli to become more salient, to stand out, to evoke deeper processing and better memory. We have not yet carried out experiments that would provide the most straightforward tests for emotional effects in selective observation.

We do have demonstrations of selective learning of happy versus sad material in a narrative by subjects who were happy or sad at the time they read it. One story was about two college men getting together and playing a friendly game of tennis. Andre is happy—everything is going well for him—whereas Jack is sad—nothing is going well for him. The events of the two men's lives and their emotional reactions are vividly described in the story, which is a balanced, third-person account. After reading, our subjects were asked who they thought was the central character and who they identified with. We found that readers identified more with the character who was in their mood, thought the story was about him, and thought the story contained more statements about him.

The subjects recalled the text the next day while in a neutral mood, with the results shown in Fig. 13.1. Subjects recalled more facts about the character with whom they had identified. Eighty percent of the facts recalled by the sad readers

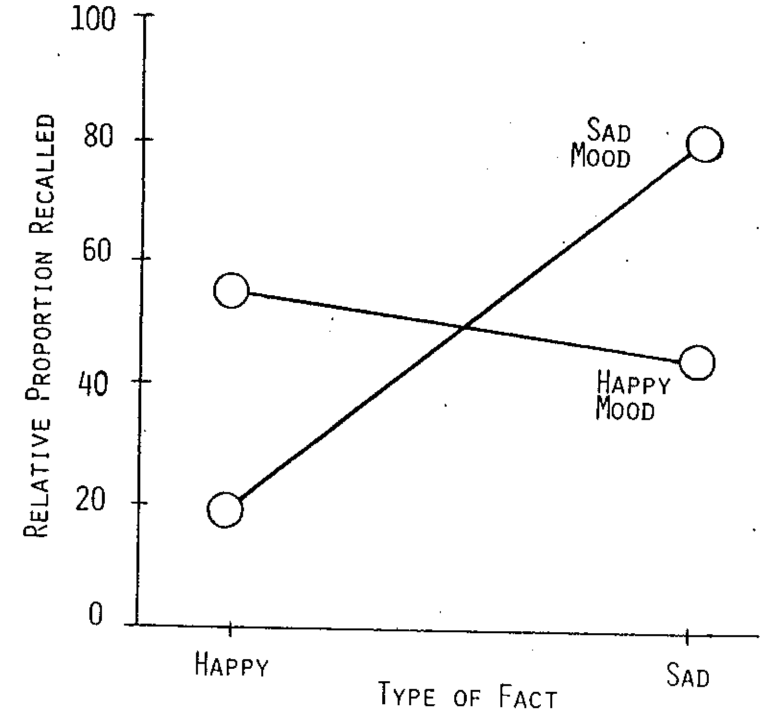

FIG. 13.1. Relative percentages of recall of facts about the happy character versus the sad character by readers who were happy or sad (from Bower, 1981).

were about the sad character; 55% of the facts recalled by the happy readers were about the happy character. This is mood-congruous selection according to the character the statement was about. It differs from state-dependent memory (to be discussed shortly), because these subjects were in a neutral mood during recall.

In another experiment (Bower, Gilligan, & Monteiro, 1981) we had subjects read some simulated psychiatric interviews, in which a psychiatrist led a patient through several sessions of hypnosis-induced age regression. The patient in the narrative briefly described a series of unrelated happy incidents and sad incidents from his life. The subjects were made to feel happy or sad while reading this by posthypnotic suggestion. Later they recalled the narrative with the results shown in Fig. 13.2.

Here again, people learned more about incidents congruent with their mood. Happy readers recalled about one and one half times as many happy incidents as sad incidents, whereas sad readers recalled about one and one third times as many sad as happy incidents. So the mood-congruity effect occurred for happy versus sad incidents related by a single character.

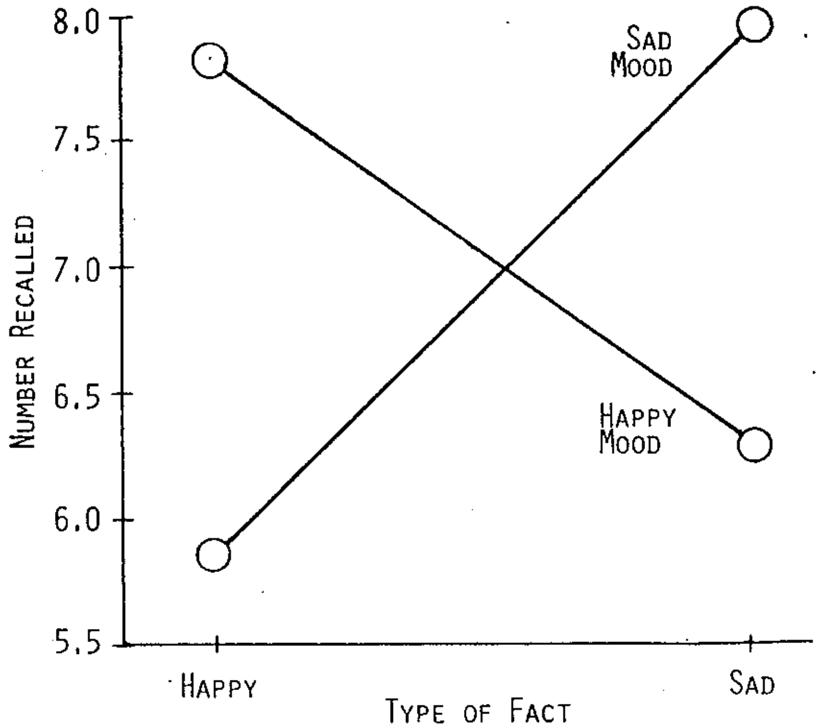

FIG. 13.2. Number of happy versus sad story incidents recalled by readers who were happy or sad. (One character in the story described both types of personal incidents, from Bower, 1981).

How are we to explain the mood-congruity effect? Without describing details here, we note that we are investigating three complementary explanations of mood-congruous learning. The first hypothesis is that subjects semantically *elaborate* more on mood-congruous material; the second is that mood-congruous narrative events are more likely to *remind* subjects of an autobiographic event; the third is that mood-congruous material causes a more *intense* emotional reaction than does incongruous material. Any or all of these immediate effects could plausibly cause greater memory for mood-congruous events. Space limitations preclude our describing this research here.

Selective Retrieval

Besides selective storage, emotion also influences which memory records are easily accessed in memory. Here, we find emotion "state-dependent" memory: People can best retrieve a memory by reinstating the emotion they were experi-

encing when they originally stored it. Thus, subjects who are happy can thereby recall better some experience stored when they were happy. The effect is not "all–or–none" but "more–or–less" recall, similar to other context changes such as altering a person's environment between learning and recalling.

State dependency may be illustrated with an early experiment by Bower, Monteiro, and Gilligan (1978). Subjects were taught to free recall two lists of 16 words, one list learned while they were happy, the other learned while they were sad. Later they recalled both lists when they were in one mood or the other. (The second entry into a mood was always achieved by remembering a different happy or sad experience than used originally.) The results are shown in Fig. 13.3. People who were tested sad recalled more of the list they had learned while sad, whereas people who were tested happy recalled more of the list they had learned while happy. Relative to control subjects who learned and recalled both lists in a single mood, subjects who learned the lists in different moods showed interference when recalling a list in its "wrong" mood, and facilitation when recalling it in its "right" mood. The results are explicable by assuming that subjects'

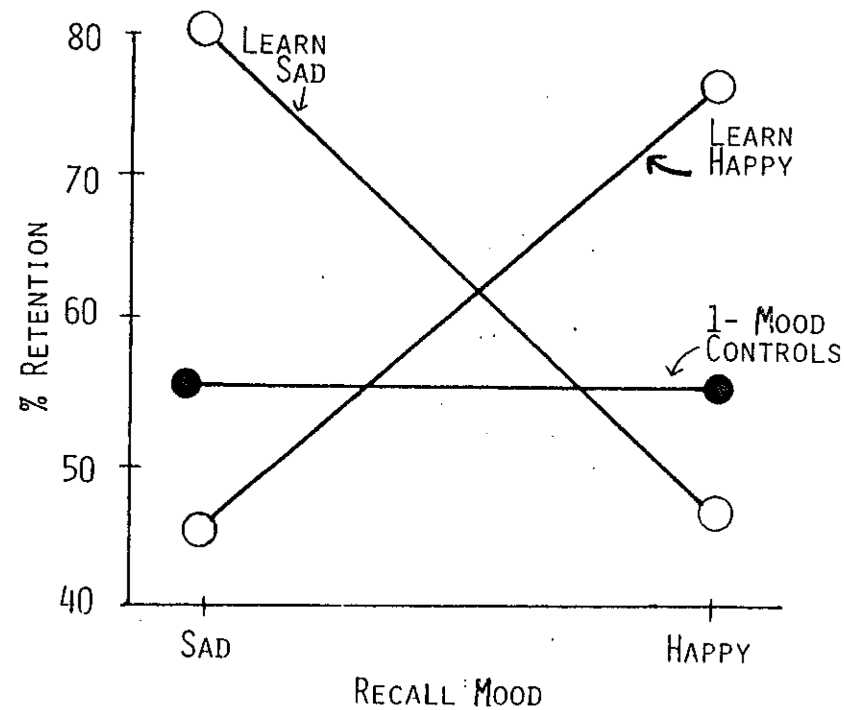

FIG. 13.3. Percentage retention scores depending on the match between learning mood and recall mood. The sloping lines refer to subjects who learned the two lists under different moods (from Bower, 1981).

mood becomes associated to the learning items, that these associations facilitate recall when the test mood matches the mood under which the targets were learned, but they interfere with recall (retrieving competing items) when the test and input moods mismatch.

This mood dependency in free recall has been replicated several times in our laboratory. We find state dependency in recall of emotional events that people have recorded in personal diaries; similarly, mood exerts a selective influence on recall of pleasant versus unpleasant childhood experiences.

State dependency does not require hypnotically induced moods, as shown by two other results. Bartlett, Burleson, and Santrock (in press) produced mood-state dependent learning with happy versus sad moods in kindergarteners and third-grade children. For inducing the moods, the experimenters simply asked the child to reimagine (and retell to the experimenter) for several minutes either a very happy or very sad experience they have had. A mood check found this was quite effective. The subjects learned two lists in different moods and were tested in one of the moods (using a different evoking experience), as in our earlier experiment. The authors reported that when happy, children recalled more of the items learned while happy (.60) than the items learned while they were sad (.49). Conversely, sad children showed the reverse pattern of recall of the two lists (sad list .58, happy list .43). Kindergarteners and third graders did not differ in the amount of state dependency shown.

A second, striking example of mood state-dependent memory was reported by Henry, Weingartner, and Murphy (1973). Psychiatric patients with bipolar manic–depressive swings were observed over several months. At several sessions throughout this period patients were asked to generate 20 free associations to each of two novel stimulus words; 4 days later patients would be asked to reproduce the same 40 words they had generated before. In addition, clinicians rated the change over the 4 days in the patient's affective state, using several mood dimensions such as degree of activation and euphoria–depression. Henry et al. found that the greater the change in patients' affective states—from mania to depression or vice versa—the more they forgot the target associations generated 4 days previously.

These several examples suggest that mood-dependent retrieval is a genuine phenomenon, whether the mood swings are created experimentally or by endogenous factors in a clinical population.

The Semantic-Network Theory

What kind of theory can explain these mood state dependent effects? Bower (1981) proposed a theory of emotion that viewed it within a semantic network model of memory similar to those common in cognitive psychology (Anderson, 1976). We excerpt the description Bower (1981) wrote before:

First, let me provide some background. Human memory can be modeled in terms of an associative network of semantic concepts and schemata that are used to describe events. An event is represented in memory by a cluster of descriptive propositions. These are recorded in memory by establishing new associative connections among instances of the concepts used in describing the event. The basic unit of thought is the proposition; the basic process of thought is activation of a proposition and its concepts. The contents of consciousness are the sensations, concepts, and propositions whose current activation level exceeds some threshold. Activation presumably spreads from one concept to another, or from one proposition to another, by associative linkages between them. A relevant analogy is an electrical network in which terminals correspond to concepts or event nodes (units), connecting wires correspond to associative relations with more or less resistance, and electrical energy corresponds to activation that is injected into one or more nodes (units) in the network. Activation of a node can be accomplished either by presentation of the corresponding stimulus pattern or by prior activation of an associated thought.

To illustrate, a simple event such as "Mary kissed me at a specific time and place" would be recorded in memory, as shown in Fig. 13.4, in terms of new labeled linkages between my prior concepts of Mary, myself, and kissing. The links are labeled S to denote the subject and P the predicate of the proposition. Learning consists of establishing these associations and increasing their strength.

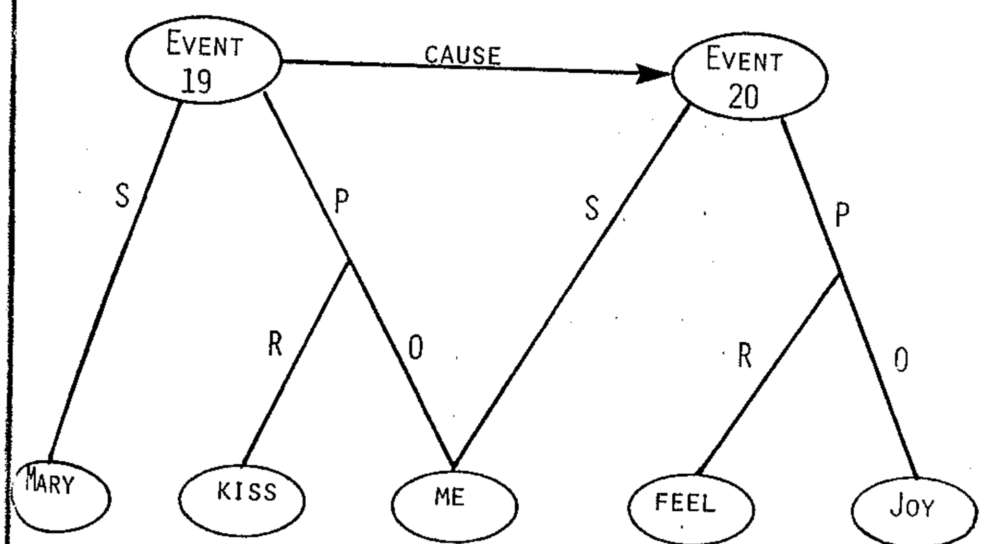

FIG. 13.4. A semantic-network encoding of a proposition ("Mary kissed me") and an emotion it causes. Lower circles, or nodes, represent preexisting concepts, and lines represent new associations, S = subject; P = predicate; R = relation; and O = object (from Bower, 1981).

Later when asked, "What did Mary do?" activation of the Mary concept will transmit activation to the Event 19 node and thence to its branches, causing the model to retrieve the other links and thus recall that "Mary kissed me."

Figure 13.4 also shows a causal link from the event to an emotional reaction—namely, the emotion of joy. The emotional interpretation of an event that creates such links is itself largely molded by cultural or personal rules of intricate subtlety—but this is not the place to elaborate on that (author's note: see the last part of this chapter). The network encodes the fact that Event 19 caused Event 20, and the latter involves a primitive node for the emotion of joy.

The semantic-network approach supposes that each distinct emotion such as joy, depression, or fear has a specific node or unit in memory that collects together many other aspects of the emotion that are connected to it by associative pointers. In a recent paper, Clark and Isen (in press) have proposed a similar conception. Figure 13.5 shows a schematic for a small fragment of the many connections to a emotion node—say, sadness for Emotion 3. Collected around this emotion node are its associated autonomic reactions, standard role and expressive behaviors (that is, the way we display sadness), and descriptions of standard evocative situations which when appraised lead to sadness. Also included are the verbal labels commonly assigned to this emotion such as sadness, depression, and the blues. Some of these various linkages are innate, while others are learned and elaborated through acculturation.

In addition, each emotion unit is also linked with propositions describing events from one's life during which that emotion was aroused. This was illustrated in Fig. 13.4 with the Mary-kissing incident that caused joy. The emotion aroused at that

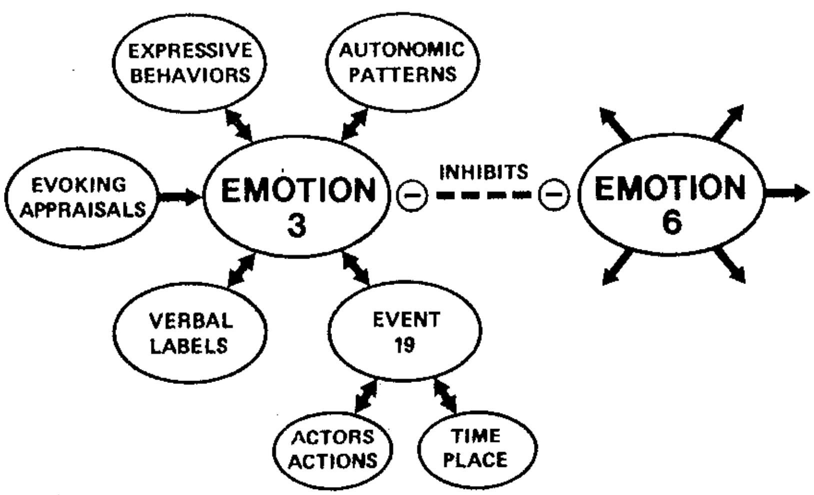

FIG. 13.5. Small fragment of the connections surrounding a specific emotion node or unit. Bidirectional arrows refer to mutual exchange of activation between nodes. An inhibitory pathway from Emotion 3 to Emotion 6 is also shown (from Bower, 1981).

time became associated by contiguity and causal belongingness with the evoking event. As a second example, the grief felt at the funeral of a friend would be associated with a node containing a description of things noticed at the funeral.

These emotion nodes can be activated by many stimuli—by physiological or symbolic verbal means. When activated above a threshold, the emotion unit transmits excitation to those nodes that produce the pattern of autonomic arousal and expressive behavior commonly assigned to that emotion. Each emotion may reciprocally inhibit an emotion of opposing quality, as fear inhibits joy and sexual arousal. If two emotion nodes are activated at once and they are not mutually inhibiting, then the subjective impression and expressive behavior pattern may be a blend or mixture of the two pure patterns. For example, sadness mixed with surprise may blend into disappointment (see Plutchik, 1980a,b).

Activation of an emotion node also spreads activation throughout the memory structures to which it is connected, creating subthreshold excitation at those event nodes. Thus, a weak cue that partially describes an event, such as "kindergarten days," may combine with activation from an emotion unit to raise the total activation of a relevant memory above a threshold of consciousness. Thus, the sad person becomes conscious of thinking about, and will recall some sad event from, his or her kindergarten days. This recall constitutes reactivation of a sad memory and sends feedback excitation to the sadness node, which will maintain activation of that emotion and thus influence later memories.

Network Theory of State-dependent Retrieval

The network view of emotional behavior has several implications. Of immediate interest is that it explains mood-state-dependent retrieval. The relevant associative connections for this process are isolated and emphasized in Fig. 13.6, which illustrates a part of the associative network encoding the learning of materials used in one of our experiments. The subject learned brief adjective–noun phrases such as "dying dog," "lost money," and "happy days" in the context of learning List 1 while experiencing Emotion 1 (see Gilligan & Bower, 1981). During learning the phrases became associated to the unit describing Context 1 and to the emotion unit active at that time. Later, when asked to recall events that occurred in Context 1, the subject activates the Context 1 node in memory, and activation spreads out from it as the subject searches for relevant items. But Context 1 is a weak, overloaded cue because it is associated with many things, so any one connection is subject to heavy interference. Suppose, however, that during recall the subject is returned to the same distinctive emotional state he or she was in during learning. Then, activation from that Emotion 1 node also will spread along its associative links to the target items, where it will summate with the activation spreading from the Context 1 cue to the items. The summation of activation at the intersection nodes causes the target events to become more accessible to recall when testing occurs in the same mood as learning. In contrast, if the mood is altered between learning and recalling, say to Emotion 2, recall suffers because the benefits of intersection from two search cues are absent; moreover, the search from the different emotion node will call up interfering associations that will compete with recall of the correct target items.

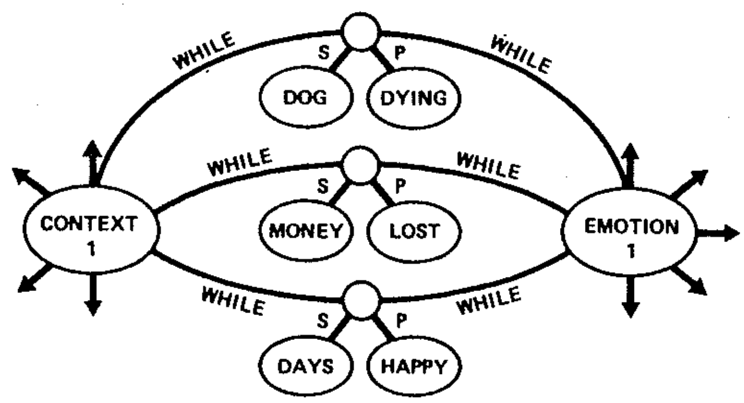

FIG. 13.6. The crucial connections for explaining mood-state-dependent retrieval. The subject has studied many adjective–noun phrases (dying dog, lost money, happy days, etc.) in Context 1 while feeling Emotion 1. The associations indicated (and many others) are weakly formed (from Bower, 1981).

This theory works more or less the same for recall of experimental items or of real-life eipsodes from the recent past or from one's childhood. The specification of the material to be recalled is partly contained within the Context 1 node as either "List A learned today" or "last week's events" or "childhood incidents."

Network theory also explains a significant qualification about state-dependent learning. This qualification is that the state-dependent effect occurs best with free recall, when minimal cues are given for retrieval of the targets, but the effect is greatly reduced when memory is tested with more adequate cues such as occur in strongly cued recall or recognition tests. To illustrate, if recall were prompted with a predicate of a phrase, such as "dying" for recall of "dying dog," recall would be much higher than free recall, and the state-dependent effect would be reduced. We have found this result, and so have Jim Eich and his associates (Eich, 1980; Eich, Weingartner, Stillman, & Gillin, 1975) in several studies with drugs. Figure 13.6 suggests why this is expected: A strong cue like the predicate of a phrase—say, "dying"—will retrieve the stored event relatively directly, because that cue is close to the target noun (of "dog") and that pathway has no competing associations. With such a cue, the memory-search process starting from the weak Context 1 cue, which otherwise would have to occur in free recall becomes unnecessary, and it was that search that we facilitated by reinstating the emotional mood of learning. Thus, the search clue provided by mood reinstatement is unnecessary when adequate retrieval cues are provided for the memory targets, so the matching of moods between input and output has relatively little effect [p. 134–136].

Face Recognition Memory

Eich's generalization that state dependency is eliminated by recognition memory tests is indeed impressive. However, as professional skeptics, we felt called upon to put it to a severe test. Consequently, Randy Gellerman and Bower searched for state-dependent memory using recognition memory for faces. We thought that of many possible stimuli the human face would be most likely to be subject to emotional biases in encoding. For example, Schiffenbauer (1970) found that, compared to neutral-mood controls, subjects induced to feel disgust or anger (from listening to a tape recording) described a set of facial photographs using many more negative than positive emotional labels. In our experiment subjects were induced to feel happy or angry as the background emotion while they studied and were tested for recognition of slides of men's faces. The slides were taken from three high-school yearbooks (supplied by Tony Bower) and were of 160 male Caucasian twelfth graders, excluding faces with highly discriminable features such as long hair, eye glasses, scars, or facial hair. Subjects were placed in a happy or angry mood, were presented 40 slides to study at a 3-second rate, then were shifted to the opposite mood and presented with a second set of 40 slides for study. Ten minutes later, subjects were returned to a happy or sad mood, took half (80) the tests, then returned to the alternate mood for the remaining tests (80). Half the test slides in each test block were distractors, a quarter had been represented earlier while the subject was angry, a quarter had been presented before when the subject was happy. Each input slide was tested either in the same or its opposite mood according to a 2 × 2 design.

If a person's facial expression is an important part of the total encoding of an event, as Zajonc (this volume) suggests, then recognition memory for a picture should suffer when the subject's test mood (and expression) differs from the learning mood for that picture. Conversely, if Eich's generalization is right, then we should find no mood-dependent memory effects on recognition memory.

The percentage correct recognition responses are shown in Table 13.1. These results showed no differences whatsoever in recognition memory according to the input mood or the output mood or their interaction. Thus, contrary to the intuition that face encoding would differ according to the mood of the perceiver, we confirmed Eich's generalization: There is no state-dependent retrieval effect

TABLE 13.1
Percentage Correct Recognitions of Faces According to Their Input Condition and Test Condition

		Test Mood	
		Happy	*Angry*
Learning Mood	Happy	64	65
	Angry	59	57
	*Distractor	83	85

with recognition measures of memory. Thus, this prediction of the model in Fig. 13.6 is confirmed.

Emotional Influences on Thinking and Judgment

The network theory expects several cognitive effects when emotions are strongly aroused. Concepts, words, themata, and inference rules connected to the aroused emotion should be primed and more available. We have examined a few instances of how emotion influences thought processes, concentrating more on fantasies or reveries rather than goal-directed problem solving.

Free Associations. One line of work examined verbal free associations.. Subjects who were happy or angry generated chains of free associations to neutral words, which often reflected their mood. The network theory implies this by supposing that the prevailing mood acts as a constant source of activation so that the most activated word associations lie on intersection points in the network between the mood and the stimulus word. Thus, the associates typically satisfy joint constraints suggested by the mood and the stimulus words.

Imaginative Fantasies. Because concepts related to the prevailing mood are primed into readiness, mood should bias the interpretation of ambiguous scenes. Because of this, and because mood also affects associations, one expects mood to influence imaginative fantasies. We have investigated how happy, sad, or angry subjects tell stories about what was happening in scenes depicted in cards of the Thematic Apperception Test (TAT). Using judges' ratings of content, we have found that the content of TAT stories frequently reflects the subject's mood; that is, happy subjects usually tell happy stories, angry subjects tell angry stories, and so on.

Snap Judgments. The network theory implies that mood will influence "snap" judgments about familiar people or objects about which the subject has stored heterogeneous impressions. Isen, Shalker, Clark, and Karp (1978) found that subjects who were feeling good after receiving a small gift were more positive in rating the performance of their TV's and their cars for a mock consumer survey. Following that lead, we had our subjects give "thumbnail personality sketches" of familiar people in their lives (e.g., cousin, uncle, teacher); some characters were described while the subjects were happy and others while they were angry. We found that their judgments were strongly influenced by their passing mood. Angry judges are merciless, faultfinding; happy judges are charitable, loving, generous.

Assuming heterogeneous impressions have been stored about familiar persons, we may suppose that one's current mood causes retrieval of primarily positive or primarily negative memories and opinions of a familiar person. In this way, the summary evaluation is biased by the availability of the positive versus negative features that come to mind. This is just a state-dependent memory effect in disguise.

Probability Estimates of Future Events. A cornerstone of rational behavior is the idea that one chooses an optimizing action by combining the utilities of prospective outcomes with their subjective probabilities. Rationality requires that the probability estimates be unbiased, or at least as objective as possible. In collaboration with William Wright, Bower has investigated how happy versus depressed moods influence subjective probabilities of future events. Subjects read two 12-item questionnaries comprised of possible future events and were asked to estimate the "objective probability" of each event on a 0-to-100 scale. Half the future events were blessings, half were disasters; within each set, half referred to personal events, and half to national or world affairs. Sample items asked the subjects to estimate the probability that within the next 3 years they would take a vacation in Europe, that in the next 5 years they would be in a serious automobile accident, and that within the next 10 years there would be a major disaster at a nuclear power plant in California. Our hypnotized subjects were placed into a happy or depressed mood, filled out one questionnaire, then were placed in the opposite mood, and completed the second questionnaire. The order of the moods and the questionnaires was counterbalanced over subjects. Subjects were asked to try to be objective as possible in their estimates, and all sincerely believed their estimates were factually based and not influenced by their moods.

Nonetheless, people's mood dramatically influenced their subjective probability estimates compared to estimates given by control subjects in a neutral mood (see Table 13.2). When happy, subjects elevated their probability estimates of positive future events and reduced estimates of negative future events; on the other hand, depressed subjects did just the reverse, increasing their probability estimates of catastrophes and lowering their estimates of blessings. These mean differences were present on 22 of the 24 event questions, despite wide scatter of the estimates and small numbers of subjects in each comparison. Here,

TABLE 13.2
Average Estimated Probabilities over Many Positive and Negative Events According to Estimator's Mood

Mood	*Event Type*	
	Positive	*Negative*
Happy	.52	.37
Neutral	.44	.43
Depressed	.38	.52

then, is the optimism of the happy person and the pessimism of the depressed person.

To explain this result, one must first note that the methods people use to make probability estimates are diverse and domain specific. A predominate strategy is to use the event in the question (say, about auto accidents) to try to remember either relevant media news programs or specific autobiographic episodes that involve such events or similar ones—accidents you have had or your friends have had. A relative frequency estimate would then be made from this available sample and extrapolated as one's future estimate. Because of mood state-dependent retrieval, memories of mood-congruent episodes would appear to be very available, hence probability estimates would be drastically biased in the expected direction (see Tversky and Kahneman, 1973). Other estimating strategies (e.g., building causal scenarios) would seem also to be similarly influenced by mood. Subjects' introspective reports were consistent with this analysis.

Although we prefer to interpret the preceding four experiments as showing the *automatic* influence of emotion upon thinking, an obvious alternative is that the hypnotic subjects are simply complying with what are the "emotional-role behaviors" implicitly demanded by the mood induction; that is, angry mood induction may be just a suggestion to play the role of an angry person in all the tasks. Elsewhere, Bower (1981) has marshalled the arguments and evidence against this demand–compliance interpretation of the results. Space limitations prevent our recounting them here.

Social Judgments. We believe that judgments about interpersonal actions are largely subjective and can be distorted by the mood of the judge or perceiver. Social acts are often ambiguous, and perceivers have to read the intentions hidden behind people's words and actions. In that reading, the emotional premise from which they begin strongly influences how they interpret behavior. Is this student sticking to his point of view with admirable persistence, or is he being pigheaded? Is this soldier's action boldly daring or is it reckless? Clearly, the judgments we make depend on how the event impacts on our goals and how we feel about the actor. Happy people tend to be charitable, loving, positive in their interpretation of others. Depressed people are quick to notice any signs of flagging friendship, to exaggerate the slightest criticism, to overinterpret remarks as personal, denigrating, and pitying. Angry people have a chip on their shoulder, tend to be uncharitable, ready to find fault, to take offense. They may "take out their anger" on an innocent bystander who had nothing to do with arousing their ire; this is the basis for "scapegoating."

These emotional influences on judgments apply just as well when people are judging themselves and their own behavior. For instance, psychiatrically depressed patients are notorious for castigating themselves for what they perceive as their incompetent, despicable actions.

Joseph Forgas, Susan Krantz, and Bower investigated whether emotion could influence subjects' moment-by-moment perception of their own behavior. Specifically, we asked whether normal college students would naturally *see* themselves as incompetent and negative if they looked at themselves while feeling socially rejected. Conversely, we wanted to see whether people would *see* themselves as behaving with positive prosocial actions if they looked at themselves while feeling happy. This required a 2-day experiment. On the first day, pairs of subjects were interviewed together for about 20 minutes about personal topics, and this was videotaped with their knowledge and consent. The next day they returned and learned how to score videotaped interviews for prosocial, positive or antisocial, negative conversational behaviors. Examples of positive behaviors were smiling, leaning forward, contributing friendly remarks; negative behaviors were frowning, looking away, grimacing, and so on. Subjects learned to score such behaviors every 10 seconds while watching two persons in a video. Following this, subjects were hypnotized. Half of them were asked to remember and replay in imagination a moment of social success when they performed spectacularly well and felt good about themselves. The remaining subjects were asked to recall and replay a moment of social failure, when they had felt embarrassed and socially rejected because of something awkward or shameful they had

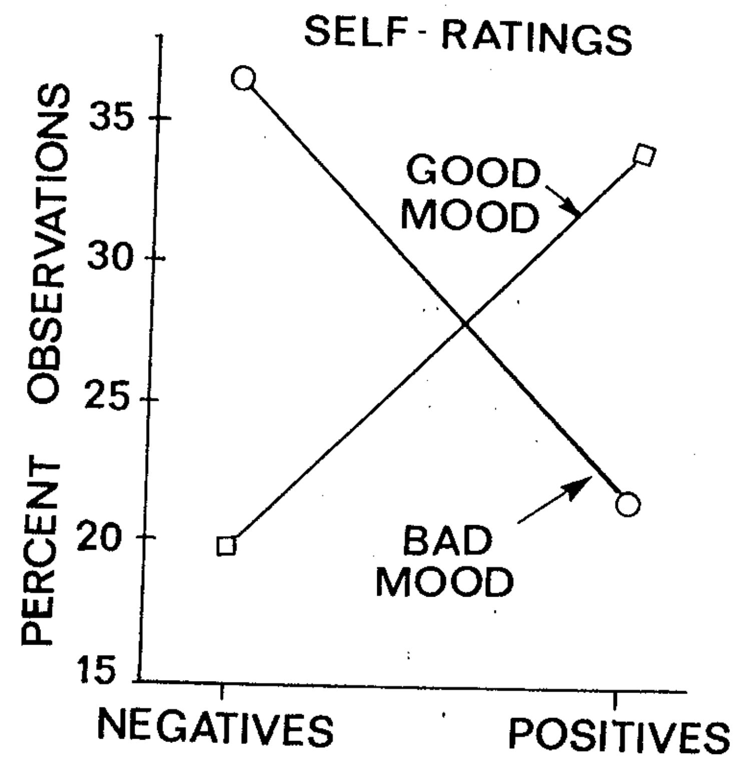

FIG. 13.7. Self-observations of positive and negative behaviors by subjects feeling good or bad about themselves.

done. Subjects were asked to maintain their mood while they looked at the 20-minute videotape of themselves being interviewed together with their partner. Every 10 seconds they were to mark at least one positive or negative behavior they observed in themselves or in their partner. Note that these are almost moment-by-moment perceptual judgments, not retrospective evaluations from memory.

The primary results are shown in Fig. 13.7 for self-ratings. Because different subjects gave varying numbers of marks, Fig. 13.7 depicts the percentages of all so-called "observations" that fall into various categories. The results show that people who felt socially rejected perceived themselves (by their "objective counts") in the videotaped interview as exhibiting more negative, socially inept acts than positively skilled acts. In contrast, subjects in a good mood perceived more positive, prosocial actions than negative actions in themselves. These differences are all "in the eye of the beholder"; neutral judges rated the two groups of subjects as exhibiting roughly equal proportions of positive and negative behaviors. The results illustrate just how ambiguous "body language" is,

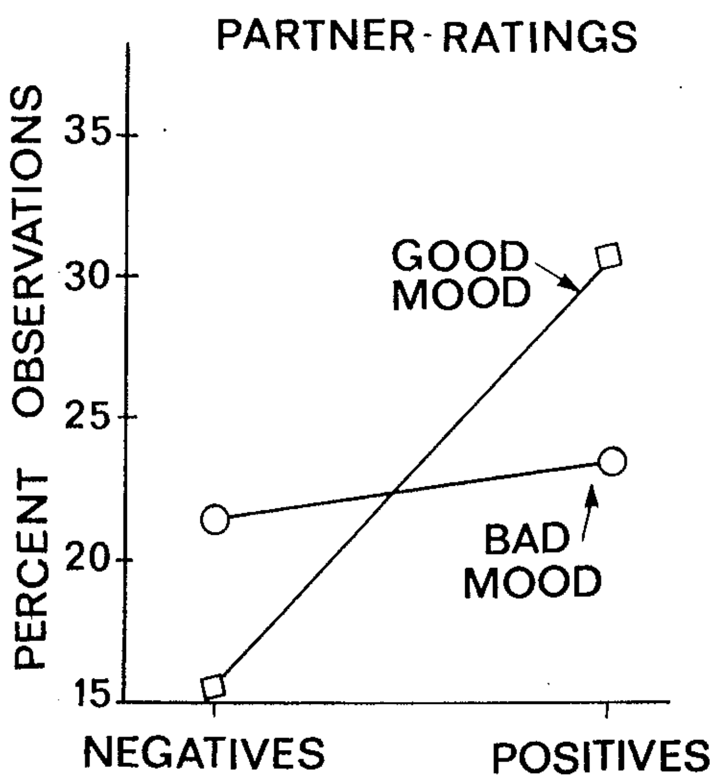

FIG. 13.8. Observations of the partner's positive and negative behaviors by subjects feeling good or bad about themselves.

because, as noted, these are moment-by-moment perceptual judgments, not retrospective evaluations. It appears that social behavior is almost a blank canvas onto which perceivers project a picture according to their moods.

The subjects' perception of their partner's behavior is shown in Fig. 13.8. Happy people see more positive, prosocial actions than negative actions being performed by their partner. Socially rejected subjects see fewer positive actions in their partner, but still their positive attributions exceed their negative ones. Thus, people who are feeling socially rejected were inclined to castigate themselves but not others; those who were feeling the glow of social success had a Pollyanna bias to see the good qualities in others as well as in themselves.

Such emotional influences on social judgments can be explained, at least roughly, by the semantic network theory. The perceiver's mood activates and primes into readiness mood-congruent concepts, hypotheses, and inference rules; and these are used in expectation-driven (top-down) processing to classify and assimilate ambiguous interpersonal actions.

To summarize these last few findings, we have shown emotional influences on several kinds of judments and thinking. This includes free associations, TAT stories, snap judgments of personality, probability estimates of future events, judgments of others and of one's own actions. We now consider an incompleteness of the network theory.

Extending the Network Theory of Affect

Parsimony recommends the idea that emotional reactions to experiences should be stored along with nonemotional features in the same memory medium, according to the same storage principles, and retrieved by the same principles. Similarly, the spreading activation mechanism is attractively familiar: An aroused emotion spreads out activation, priming into readiness related concepts, words, themata, and inference rules. This enables the model to react differently to the same situation depending on its current emotional state. Further, the idea that an emotion becomes associated to coincident events and can later act like a retrieval cue suffices to explain mood state-dependent memory.

Whereas the network theory deals with the cognitive consequences of an emotion after it has been aroused, its major incompleteness is that it does not address how emotions are aroused in the first place. Figure 13.6, for instance, depicts "evoking appraisals" that feed excitation into the emotion node, but that part of the system was unanalyzed. That was because in earlier work Bower and his associates were interested in the cognitive *consequences* of emotional states, not their antecedent *causes*. Nonetheless, we feel that the "pattern-recognizing" aspects of emotional arousal should be addressed, and so we propose the "blackboard" control structure described following.

A theory about cognitive consequences of emotion needs to deal with emotional appraisals for several reasons. A first reason is that one major consequence of emotion is to change the way people appraise or evaluate situations. This was our point about mood affecting social judgment. To explain such results, then, one must model the appraisal process.

A second reason to deal with appraisal is to enable the model to assess properly its emotional attitude toward people, objects, and issues. Consider an illustration: We witness an argument that culminates when a bully, Sam, beats up a smaller child, Johnny. The model will store a description of that event in memory linked to the emotion it felt at the time, such as anger. Later we ask the memory model how it feels about Johnny. Can the system calculate how it feels about a person or entity by simply averaging the evaluations linked to episodes in which that entity occurs? This "inheritance of affects" strategy gives the wrong answer in this case, because the model should not feel angry at Johnny. The model needs a more discriminating appraisal system in the first place, one that will assign praise or blame appropriately to the participants in an emotional episode. Thus, an appropriate model would store anger with the bully (Sam), sympathy with the victim (Johnny), and perhaps distress at the episode of Sam's hitting Johnny. But that "assignment of responsibility" uses capabilities beyond the initial network theory. Our blackboard appraisal approach can deal with such examples.

A third reason for elaborating the control structure that uses the network data base is to distinguish "hot" versus "cold" uses of emotion concepts. The former network approach basically had one node for each emotion, say "fear," and any reference to or use of that concept would somehow "turn on" that fear node in the network. But that is also the operation the model identified as "feeling afraid" subjectively. So that leads to the absurd implication that people always feel afraid when they refer to the concept. But surely, people discourse cooly about emotions many times without feeling the emotion at all. Consequently, we should distinguish the node corresponding to the concept of fear versus the node for experiencing fear itself. But then the model needs a control structure and a pattern recognizer to decide which node to activate in different circumstances.

Fourth, a control structure and appraisal system is needed to deal with the phenomena of hot versus cold *replaying* of an emotional memory, and *reinterpretation* or reappraisal of earlier emotional memories. Consider the replaying of memories in imagination. We are impressed by the fact that for many emotional memories this can be done either *with* feelings or *without* feelings. How does a model of the mind account for the difference in these cases? A plausible conjecture is that the feelings arise as you unpack, develop, and recreate your thoughts at the original experience rather than skimming over an abstract summary of the episode. Our blackboard model deals with this by supposing that to reexperience the emotion of an event the system must recreate as closely as possible from memory the state of the "blackboard" at that time, then appraise that "blackboard situation" again as though it were fresh, and then "turn on" the appropriate emotion nodes. But to do this, one must have a mechanism for appraisals.

Reinterpretation should perhaps better be called reappraisal. People often review earlier experiences and modify their emotional appraisal of the events. Thus, a spouse's remark that upset us at the time may be reappraised later as an innocent remark designed to comfort us; a stream of pleasing flattery may be later reappraised as an insincere manipulation to benefit the flatterer. We change emotional appraisals for several reasons: First, different facts may come to light that alter *cognitions* about some person or event (e.g., facts that prove their innocence); second, we may change our basic *values* or attitudes (e.g., toward marijuana smoking) between the original experience and the reappraisal; third, we may change the *importance* assigned to the event because our goals or priorities change; fourth, the passage of the event replaces our anxious *uncertainty* with calm hindsight about its outcome; fifth, our *mood* may change from, say, anger to charitable benevolence, so in reappraisals this new, charitable light is cast upon events. These five causes of reappraisal—changes in evidence, in values, in importance, in uncertainty, in moods—are discussed later within our expanded blackboard model. The key idea is that the reasons for an initial emotional appraisal of an event are stored in memory along with the event and its appraisal. This data structure permits the model to reexperience the event by replaying one's original emotional reaction to it, or to replay the event and reinterpret it using modified appraisal rules, perhaps even incorporating new facts into the second appraisal. This new appraisal may then be linked as an update to the old event structure in memory. This process then permits the model to say things like "I used to like Bill's flattery. But then I learned that he was an insincere manipulator, so now I feel ashamed that I was duped."

Having listed some incompletenesses of the former network theory of emotion, we now introduce a general framework for modeling emotional appraisals.

The Blackboard Control Structure

We believe that a model of how people arrive at emotional appraisals must be complex because it requires integration of several knowledge sources. Consequently, any model will require a control structure that facilitates the coordination of many knowledge sources. In artificial ingelligence programs like HEARSAY II and CRYSALIS, this coordination is carried out using a "blackboard" control structure. The blackboard control structure is well-suited for interpretations that develop as multiple knowledge sources contribute their advice opportunistically and asynchronously. The blackboard model develops interpretations of a situation incrementally; interpretations gain support from the several knowledge sources. The partial interpretations produced by knowledge sources are called hypothesis elements. Typically, the input to a knowledge source is the output of another knowledge source; for example, a situation might have to be identified cognitively by one expert before its emotional significance can be appraised by other experts.

The blackboard is also a model of the focus of attention; the model's attention is focused on the hypotheses on the blackboard. Several competing hypotheses may be under consideration simultaneously. However, attention eventually focuses on that hypothesis for which the evidence is greatest; thereafter, the alternative hypotheses slowly fade from the blackboard. In this respect, the blackboard has many of the properties of short-term memory or working memory. In fact, no misunderstanding will arise if psychologists think of the blackboard as their construct of working memory.

We think of the contents of the blackboard as separated for the moment from the knowledge sources that use it. However, in a system like ACT (Anderson, 1976), working memory is just the activated portion of the long-term memory network plus some activated rules. To recast the blackboard model within ACT would require that the "blackboard" be considered to be those long-term memory nodes (concepts or propositions) whose activation level exceeded a certain threshold. We do not attempt this translation here. ACT and the blackboard model may be notational variants of one another, but we find the "blackboard" level of description easier to follow.

Emotional Knowledge Sources

We believe that people interpret a situation cognitively and then appraise that cognitive interpretation in arriving at the type and level of emotional reaction they will have. Thus, any model must begin with cognitive "pattern-recognition" knowledge, which enables it to recognize objects, people, events, and so on. This "cognitive interpretation" (C-I) knowledge is applied to material on the blackboard to arrive at internal symbolic representations of the external environment. We model C-I knowledge in terms of interlocking sets of productions, as in HEARSAY-II. Cognitive interpretation is shown as the early phases diagrammed in Fig. 13.9. Cognitive interpretation itself is a multistage process moving from sensory–feature extraction to conceptual meaning interpretation. The cognitive interpretation may vary in its level of refinement and detail as more processing occurs, and an emotional response can be triggered at any level or stage in processing.

The second kind of knowledge assigns an emotional appraisal to the cognitive interpretations; we call this "emotional interpretation" or E-I knowledge. We model this also in terms of productions called E-I rules. Because this is the central component of our model, we postpone its discussion for the moment.

Third, Fig. 13.9 shows that the intensity of the emotional reaction posted on the blackboard is adjusted depending on the personal importance of the event. Importance assignment requires knowledge of the person's goals and basic values. Once goals are known, then a partial ordering of importance of subgoals can be constructed. As a further factor, intensity of an emotional reaction may be adjusted by the uncertainty of the event. Unexpected events usually generate

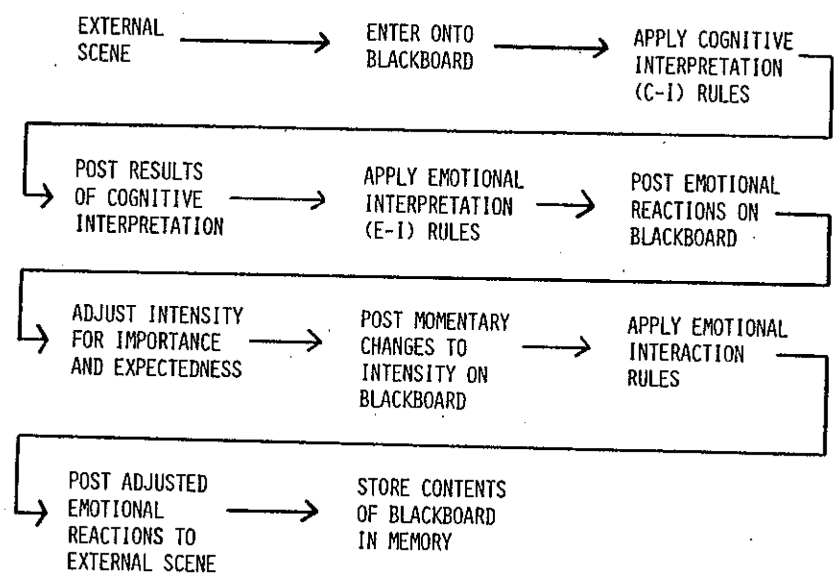

FIG. 13.9. Changes to the state of the blackboard result from knowledge sources posting hypotheses during appraisal of an emotional situation.

stronger emotional experiences than expected ones. We say little about either of these knowledge experts except to note that they are clearly needed by the model.

Fourth, Fig. 13.9 indicates that the emotional reaction is modulated by its interactions with other emotions that are either ongoing or evoked by the event in question. Thus, a sad or frightening event will decrease one's happiness, anger may intensify one's depression, and so on. For convenience, we formalize these interactions by rules that adjust an ongoing emotion in light of a temporary emotional change created by an incoming event. The rules in these cases could just as well have been implemented as inhibitory or excitatory connections among emotion nodes, like a system of lateral inhibition. One's choice on this implementation issue has certain implications. For instance, explicit rules about emotional interactions provide declarative knowledge that the model can refer to and reason about (e.g., to manipulate others' emotional reactions). Conversely, an inhibitory link between two emotion nodes is like tacit procedural knowledge that "does its job" but cannot be reasoned about or altered. We prefer the rule formulation because the model requires flexibility to change interpretations due to experience.

Emotional-Interpretation Rules. The most important part of the appraisal system is the emotion interpretation knowledge, which we formalize as many E-I rules or productions. These specify for a given cognitive interpretation an appropriate emotional appraisal. Each rule has a left-hand side (LHS) that represents a cognitive interpretation, and a right-hand side (RHS) that specifies an alteration of one or more emotions. Such E-I rules may have sophisticated conditions and outcomes. Whereas highly specific E-I rules describe individual events (such as a bully hitting a victim) in their LHS's, we believe that as some E-I rules are generalized, their LHS's come to specify goals and motives. We believe that there is a class of events, called "goal events," that includes goal achievement, goal blocking or frustration, goal abandonment, plan interruption, goal conflict, and so on. These goal events are present in the LHS's of general E-I rules, and emotion results when they are instantiated. We believe that individuals maintain a large set of background goals, such as "preserve well-being," as well as more specific goals activated as they work on a specific problem. One's goals may perhaps be organized into hierarchies along the lines described by Carbonnell (1979). The importance of goals is that they direct processing in the system. When goals are achieved or thwarted E-I rules should turn on the appropriate emotion. For instance, if while pursuing a top-level goal you notice that a required subgoal is already true in the world, then you should feel happy. An appropriate production is:

> IF a subgoal is discovered to be true,
> THEN increase happiness.

Similarly, if a goal is frustrated, the emotion of frustration or anger results.

The left-hand side of such rules can be generalized by replacing constants by variables or by deleting restrictive clauses. Similarly, a rule may be specialized or refined by the reverse operations, adding clauses or specifying a variable to be a particular constant. For instance, a more specific version of the earlier rule would be:

> IF you're starting on an auto trip and you find that your friend
> already filled your car with gasoline,
> THEN increase happiness.

A generalization of the earlier rule would be

> IF you reduce any difference between your
> current state and the goal state,
> THEN increase happiness.

The E-I rules may be hierarchically organized according to the generality of their triggering conditions, their LHS's. They can also be organized by the sophistication of cognitive interpretation required to match their LHS's. For example, an unexpected loud noise provides an immediate stimulus for which even infants have an emotional production; on the other hand, snide social remarks require sophisticated processing to determine whether one has been complimented or insulted.

The LHS's of E-I rules are partial descriptions of some state of the world, so cognitive interpretation is needed to produce this information. The knowledge involved in cognitive interpretation maps from external stimuli to an internal representation of the world; E-I rules map from these internal representations to emotional reactions. The knowledge used in cognitive interpretation tells us what is going on in the world; E-I rules tell us how we feel about it. Interestingly, cognitive interpretations can be disputed in a way that emotional interpretations cannot; one's perceptions can be denied, but not one's feelings. Emotional interpretation is subjective. Although some common, general, E-I rules exist, individuals surely differ to a large degree in their stock of E-I rules.

Most E-I rules are learned, and some are idiosyncratic. However, we expect that many primitive and adaptive behaviors can be modeled by innate E-I rules. For example, almost anyone—from infants to adults—will become upset by very loud, unexpected noises; similarly, everyone experiences frustration when a goal is thwarted. Between the idiosyncratic E-I rules and the "hardwired" ones is a range of rules that are culturally learned. For example, events that caused our grandmothers to shrink in horror, such as the wearing of skimpy bathing suits, are commonplace today, and those that were commonplace then, such as slaughtering and cleaning chickens, are repugnant to most of us today.

We think that E-I rules should not specify a complete emotional state in their RHS's, but rather an adjustment to the current level of a specified emotion, such as "increase fear," or "decrease happiness." A person's current emotional state may be described as the activation level of a set of N emotions like fear, anger, happiness, sadness, disgust, and so on. We may then conceive of any state as a $1 \times N$ vector of intensity values, varying, say, on a 0 to 1.0 scale. An E-I rule increases or decreases one or more emotions, say, by an associated linear operator.

By admitting N emotion values at differing levels, the model can deal with the phenomena of ambivalence and mixed, conflicting emotions. We assume that the central processor can report the activation level of any emotion, and momentary behavior will be selected usually by the stronger emotion. If several emotions can be satisfied by one behavior, then that behavior will be favored by the sum of their strengths. As noted, conflicting emotions interact; for example, fear inhibits joy, fear inhibits anger, etc. Thus, when the cognitive situation changes and one emotion is increased, a dynamic adjustment process may be launched as the other emotions adjust to the perturbation.

Importance, Unexpectedness, and Emotional Intensity. The RHS's of E-I rules adjust the intensity of a mood; for example, they "increase happiness" or "decrease fear." Two other factors that influence the intensity level of an emotion are importance and unexpectedness. These "scale" the adjustments to intensity specified by E-I rules. For example, consider an E-I rule that specifies an increase in intensity of anger of, say, 40% in response to seeing a bully attack a weakling. This increase should be scaled upward even further if the weakling is especially important (e.g., if he is a member of your family), and it should be scaled downward if he is unimportant (e.g., if it is a random act of violence on TV news). Similarly, the unexpectedness of an event scales the intensity of one's emotional response to it; a surprising event is assumed to adjust intensity upward, whereas an expected event would produce lower activation.

It is difficult to specify the rules that adjust intensity depending on importance, because what is important is domain specific. For example, in social relations, truthfulness (or the lack of it) is important, in chess, control of the center of the board is important, in medical diagnosis, accurate tests are important, and so on. Although there may be no domain-independent rules for assigning importances, some rules are plausible for a general domain like planning. For instance, the importance of a goal would depend on the number of subgoals it subsumes in a plan. A distinction is made between conjunctive goals that must all be achieved for a plan to work, and disjunctive goals, any of which will allow the plan to work. Conjunctive and disjunctive subgoal trees show contrasting asymmetries in the importance attached to success or failure of individual subgoals. In a disjunctive tree, the success of any subgoal increases with the number failed previously. In a conjunctive tree, the failure of any subgoal has equal importance and is to be avoided, whereas the importance of a subgoal success increases with the number of subgoals previously achieved.

Expectedness and unexpectedness also scale intensity. In general, unexpected events lead to more intense emotions. A blackboard model can judge unexpectedness because it continuously posts top-down hypotheses about the state of the world. Thus, rules for judging unexpectedness are of the form:

IF the blackboard does not contain a top-down
hypothesis for this event,
THEN the event is unexpected.

Clearly, more sophisticated rules can be formulated that deal with partial matches and degrees of surprise, but this rule illustrates the general principle.

The modulation of emotional intensity by expectedness is illustrated by our reactions to jokes and suspense stories when they are novel instead of overly familiar.

Interaction Rules. Interaction rules also modulate the intensity of emotional reactions. The LHS's of these rules contain at least one clause describing an emotional interpretation, and optional clauses describing other emotions and cognitive interpretations. An example of an interaction rule is the following:

IF intensely happy, and
an emotional interpretation of sadness is posted for
an event,
THEN decrease the intensity of ongoing happiness by 10%,
cut the increment in sadness by 50%.

The intensities of emotions that are posted are scaled for their importance as described earlier, so if the sad event were very important its intensity would be reduced less than that of happiness.

A second function of interaction rules is to specify an intuitive progression as emotions are aroused in sequence by events. For example, if one experiences fear but then discovers that there was nothing to be afraid of, he will feel relief followed perhaps by a feeling of shame or foolishness.

A third function of interaction rules is to "calm down" the emotions after instigating circumstances are removed. A simple rule for dissipating emotion is this:

IF emotion E is active, and
the emotional interpretation that evoked E is no
longer active on the blackboard,
THEN decrease the level of intensity of E by 50%.

To summarize our comments, Table 13.3 cites the knowledge sources and specifies their input–output relationships. We have discussed cognitive interpretation (C-I) rules, emotional interpretation (E-I) rules, importance and unex-

TABLE 13.3
Input–Output Relations Among Four Kinds of Knowledge Involved in an Emotional Appraisal

Rules	*Input*	*Output*
Cognitive Interpretation (C-I)	Sensory Stimuli, or RHS of C-I rules	Symbolic representation of external environment
Emotional Interpretation (E-I)	RHS of C-I rules	Adjustment of current emotion
Importance and Unexpectedness	RHS of C-I rules	Adjustment to intensity of an emotion
Interaction	RHS of E-I, and (RHS of C-I rules)	Adjustments among emotions, Adjustments to intensities of emotions

Examples

C-I: IF you see squiggles of a certain type on an EEG record,
THEN infer an ongoing petit mal seizure

E-I: IF you discover you have epilepsy,
THEN increase fear

Importance: IF A and B are disjunctive methods of achieving G, and A has a lower cost than B,
THEN A is more important than B

Interaction: IF you are happy, and a frightening stimulus occurs,
THEN decrease happiness.

pectedness rules, and interaction rules. All these interact on the blackboard to influence the intensity and quality of emotional reactions to either an external situation or replaying of a memory or imagined scene played out on the blackboard.

Emotional Situations and E-I Rules

In a chapter describing a model for recognizing emotional situations, we should try to characterize generally situations that cause emotional reactions. This is difficult because the evoking situations vary so widely. Nonetheless, some tentative remarks are in order about a taxonomy of emotional situations.

A first type of event that arouses emotion is unexpected, intense, startling stimuli. This class includes painful stimuli. Equally arousing are physical events like slipping, falling, or losing one's balance. These events will have innate E-I rules for turning on the emotions of surprise, pain, and/or distress.

A second source of emotions are cues that have become conditioned to these intense or painful events, which lead to feelings of apprehension.

The third and largest source of emotions are events surrounding the achievement, interruption, or thwarting of motivated, goal-directed activities. Evolution has equipped us with a set of homeostatic mechanisms for regulating our bodily needs, to obtain enough oxygen, food, water, benign temperatures, and so on. Deprivation arouses the relevant need state that motivates actions to reduce the need, with goal achievement accompanied by excitement, satisfaction, and pleasure. These are the prototypic "positive reinforcers" that traditional learning theory talks about.

Learning theorists such as Hobart Mowrer (1960) believed they could extend the previous two categories of positively or negatively valenced reinforcers, thus classifying any stimulus presentation as pleasantly positive or aversively negative. Emotions were then conceived as reactions to the presentation or removal of these positive or negative events, or as reactions to cues signalling the imminent presentation or removal of these events. Thus, in Mowrer's scheme, joy results from getting an ongoing positive activity leads to frustration and anger, loss of a sustained positive relationship causes sadness or anger; presentation of an aversive event leads to pain, distress, and perhaps anger; anticipation of an aversive event leads to fear or anxiety; removal of an aversive event causes relief; anticipation of that removal leads to hope.

A basic problem with this classification is that it presupposes a prior classification of events as positive or negative. But that classification surely varies according to individuals' goals and cognitions. For instance, achieving a top-level goal like losing weight may require obese people to alter their customary evaluation of eating sweets or sitting in very hot steambaths.

Another problem with Mowrer's scheme is that it misses the point that humans' emotional reaction to an event varies with their explanation of the causes of the event. This is illustrated clearly in Weiner's research reported in this volume. People have elaborate theories for explaining their positive events (successes) or negative events (failures) along the causal dimensions of internal versus external factors, controllable versus uncontrollable factors, and stable versus unstable factors. Thus, a success attributed to one's own efforts or ability leads to pride, whereas failure ascribed to lack of effort or ability leads to depression. A negative, controllable event that one attributes to others causes one to feel anger. Guilt occurs when one causes a negative outcome for a friend when it could have been prevented. One feels pity toward others who are in a negative state due to uncontrollable factors. One feels gratitude toward others when one receives positive events due to their personal control.

Weiner's taxonomy of causal factors leading to different emotions fits neatly into the E-I rules envisioned for our blackboard model. One can easily write productions whose left-hand sides are values of the several causal attributes of an event, and the right-hand side is an emotional interpretation. For instance,

> IF someone acts in a negative way
> and the factors causing the act are
> internal, stable, uncontrollable,
> THEN feel pity at $\Theta = .7$
> and feel anger at $\Theta = .2$.

> IF someone's action leads to a positive event
> for you and that outcome was under under his control,
> THEN feel gratitude at $\Theta = .6$
> and feel happiness at $\Theta = .4$.

The point is that taxonomies of social situations that generate emotions (such as Weiner's or Roseman's, 1979) can easily be translated into our model as E-I rules that turn on emotions appropriately depending on the cognitive appraisal of the social situation. Of course, such E-I rules are learned during acculturation of each individual and vary somewhat across cultures.

Having discussed briefly the range of situations that lend to emotional reactions, we now turn to illustrating the blackboard model for emotional evaluation.

An Illustration of the Blackboard Model's Operation

To illustrate the blackboard system, we walk through an example supplied by a friend; call him Fred. His apartment had been burglarized a week ago, and the police told him just today that the burglars were in fact murderers escaped from prison. The event in question occurs late that night when Fred is awakened by a noise in the hallway outside his apartment. Initially, he thinks the noise was caused by burglars, experiences fear, considers telephoning the police. But then he hears singing, infers the noise and singing is coming from his friendly neigh-

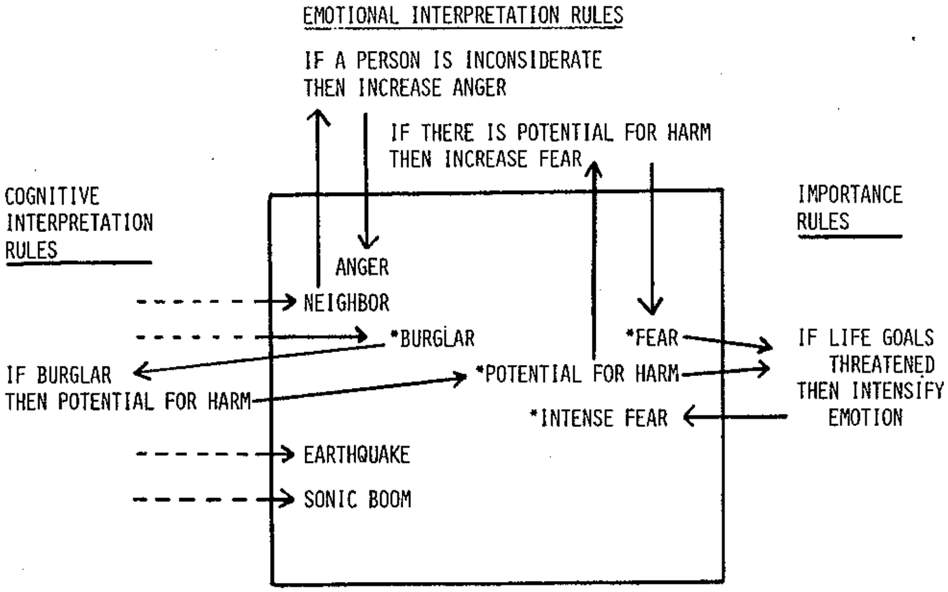

FIG. 13.10. The state of Fred's blackboard after interpreting a noise as a burglar.

bor, Bill, who's staggering home after drinking too much. Fred then feels angry at Bill for inconsiderately frightening him in the middle of the night.

Figure 13.10 shows the state of Fred's blackboard after he has just heard the noises. Cognitive interpretation rules provide several possible explanations—a burglar, a neighbor, a sonic boom, an earthquake (this is California)—and E-I rules assign an initial emotional interpretation to each. We use concept names like "burglar" as shorthand for an existential proposition like "There is a burglar out there." Some of these hypotheses have only a fleeting life. The sonic boom and earthquake hypotheses are quickly discarded when Fred localizes the noise as coming from within his building. That leaves the blackboard containing two competing hypotheses about the noise. At this point top-down expectations tilt the balance toward the burglar hypothesis. The knowledge that generates these expectations is not shown in Fig. 13.10; however, Fred's learning about the murderers that day will have primed a class of hypotheses about burglars and the fear they should provoke.

The burglar interpretation is temporarily favored over the neighbor interpretation, but both remain on the blackboard. The "importance" source assigns a large intensity increment to the fear felt on this occasion because burglars are potentially dangerous. By focusing attention on the burglar interpretation, the model now becomes biased to interpret ambiguous stimuli as evidence of a burglar. Thus, small, metallic sounds would be interpreted as the burglar trying to pick a locked door in the hallway.

At the next juncture in the episode, Fred hears singing and interprets it as Bill's drunken song. The Bill interpretation gains prominence over the burglar interpretation. In so doing, the emotional interpretation of anger gains prominence over that of fear. Although the new evidence makes the neighbor interpretation dominant, there is still a residue of the burglar interpretation, so that fear persists for a while.

Upon identifying the noise as drunken Bill, Fred becomes very angry. We explain this emotional adjustment in terms of interaction rules. These rules cause Fred to become more angry than he would have if he hadn't been frightened at first.

Selection and Application of E-I Rules

The E-I rules are organized in a generalization hierarchy according to their LHS's. This hierarchy sits in long-term memory, with each rule "looking down upon" the blackboard for its triggering conditions. Figure 13.11 shows a portion of the E-I hierarchy for interpreting and reacting to simple acts of aggression. The LHS of each E-I rule is represented as a ACT network. The notation at the terminal leaves of the rules are explained shortly.

The cognitive system produces a description of an external situation as a set of propositions such as "Sam hit Johnny" and "Sam is older and bigger than Johnny." We assume that such event descriptions are output by the perceptual parser and are entered onto the blackboard as ACT-like tree structures. (This is no substantive restriction because ACT structures are expressively equivalent to the first-order predicate calculus.) The next step is to try to match some left-hand–side of an emotional interpretation rule to this event description on the blackboard. As noted, the LHS's of E-I rules are ACT-like network structures or Boolean combinations of them, as illustrated in Fig. 13.11. When an E-I rule is instantiated, the RHS assigns emotional reactions to the event as a whole and to the participants in the event. The E-I rule itself tells the system how to assign the appropriate emotional charges. It does this by linking one or more central emotion nodes (say, anger and sympathy) with designated intensities into designated spots in the matching conceptual structure. This action is carried out by the links from the terminal nodes of the E-I structures. Once these linkages are made, activation flows over them and "turns on" the appropriate emotion node.

We think of E-I rule selection as pattern matching of LHS's of rules to the blackboard event or its immediate inferences. More precisely, we view rule selection as a three-part search process. The first part partitions the space of E-I rules into at least two categories: mood-congruent E-I rules and E-I rules that are not mood congruent. The second phase of search produces a list of all mood-congruent E-I rules that match the event on the blackboard. If this process fails to

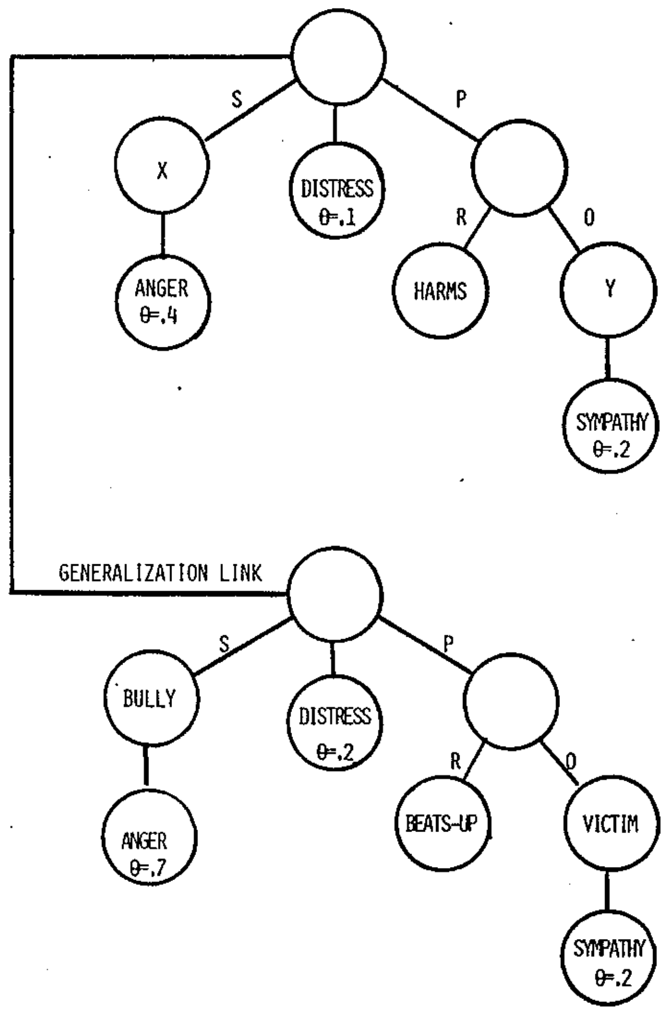

FIG. 13.11. A specific rule and its generalization, replacing some constants by variables. The modulator nodes for emotion are connected to the main concepts.

find any matching mood-congruent rules, it searches the category of nonmood-congruent rules. The third phase of search ranks the rules found in the second phase, effectively narrowing the list down to a single rule. The choice of criteria for picking one rule from several applicable rules is a strategic one. We favor picking the maximally specific E-I rule, that is, the one that binds to the greatest number of constants in the event on the blackboard. For example, if the event is "Sam hit Johnny," then the third phase of search should favor an E-I rule that specifies "Sam hit someone" over one that says "Someone hit someone"; and it should prefer both to a rule that says "X hit Y."

The first phase of the selection process is consistent with the network model of state-dependent memory; that is, it can be implemented by spreading activation in a network of E-I rules. The space of E-I rules is easily partitioned into mood-congruent and nonmood-congruent rules by spreading activation from the currently active emotion node to all rules that mention the emotion in their RHS's. This elevates the activation level of mood-congruent rules so that the system favors them over others.

The goal of the second phase of rule selection is to list those E-I rules whose LHS's match the event on the blackboard. Because this phase of rule selection considers mood-congruent matches before those that are not mood congruent, the search process will perpetuate the current mood if possible.

This matching phase is complicated by a number of factors. First, complete and literal matches are relatively rare; instead, a LHS may be a partial match to the event on the blackboard, or it may match variables to constants, or both. For example, a complete, literal match to the blackboard event "Sam hit Johnny with a stick" is simply "Sam hit Johnny with a stick"; a partial match is "Sam hit Johnny"; a matching of variables to constants is "Sam hit someone." Another reason that matching is complicated is that we intend matches to be made not only to blackboard events but also to immediate inferences from them. For example, we want to infer from the facts that "Sam hit Johnny" and "Johnny is much smaller and weaker than Sam" the conclusion that "A bully hurt a weakling." We assume that this reasoning is done by cognitive interpretation (C-I) rules, and that the conclusions are posted on the blackboard, where they are matched by the LHS's of E-I rules.

Figure 13.12 shows the structure that results from matching an E-I rule (P238) to "A bully hurt a weakling." The E-I rule makes assignments of emotions to the components of the event; it associates anger at a certain intensity with Sam, the bully, sympathy with Johnny, the weakling, and distress with the episode as a whole.

We have several comments about the data structure in Fig. 13.12. First, the emotional assignments have been determined in this instance by a general rule; the event need not have happened for the model to have predicted how it *would* feel should an event of this type arise. Second, the model is being told to feel three different emotions at once—anger, distress, and sympathy—although these have different evoking objects. Whereas the more strongly aroused affect will tend to dominate thought and action, the lesser emotions are nonetheless present and produce mixed or alternating feelings depending on the focus of the model's attention.

Third, this data structure with its emotional assignments is what will be stored in long-term memory about this episode—specifically, the event with its emotional assignments will be stored with a token of the E-I rule used in this appraisal. Later we discuss the benefits of such memory records.

Fourth, the nodes in Fig. 13.12 with emotion labels are called "modulator" nodes. They are three-way links. One link is made upon this occasion directly to

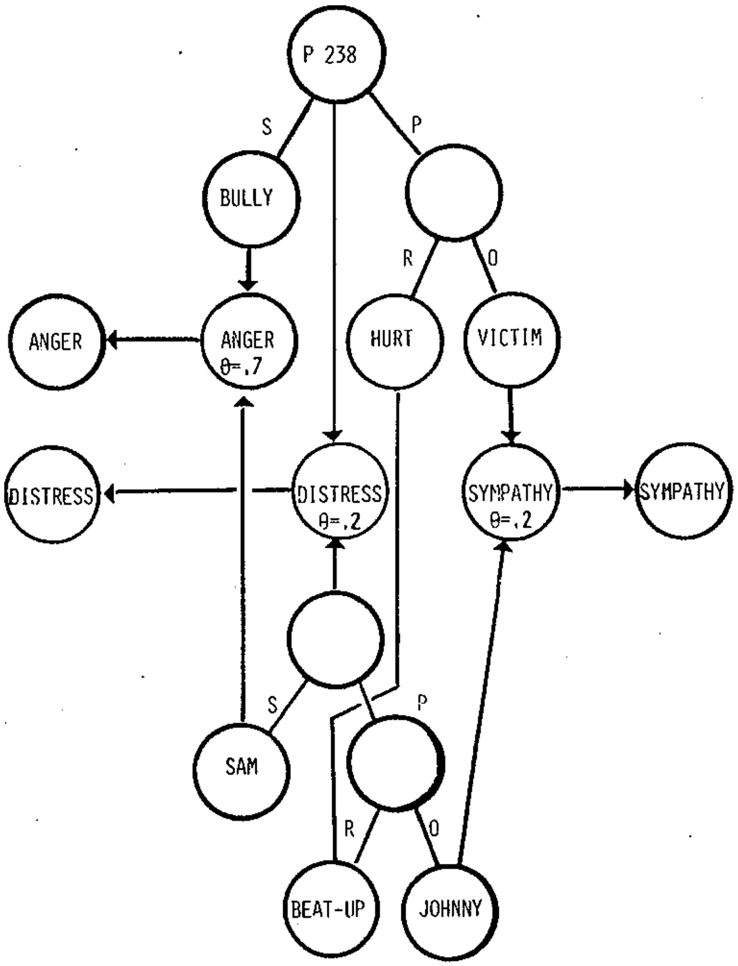

FIG. 13.12. The episode "Sam beat up Johnny" is matched to an E-I rule "Bully hurts victim." Emotional reactions to the elements of the episode are specified by modulator nodes.

the parent emotion node such as anger or sympathy. The second link goes to the abstract variable in the E-I rule that has been instantiated in the instance (for example, "bully" or "weakling"). The third link goes to the specific concepts of the event that instantiated particular roles of the E-I rule (Sam and Johnny). We think of modulator nodes as gates or valves that let activation flow from the active nodes on the blackboard (namely, Sam and bully) through the modulator gate over to the parent emotion node, thus causing that node (anger here) to be aroused at a specified level. This results in a token of that emotion being placed on the blackboard.

Fifth, the Θ values on the modulator nodes in Fig. 13.12 summarize the *intensity* of the emotional impact of the event upon each of its component concepts. We represent the intensity of activation of an emotional node on a scale from 0 to 1.0. Then the impact of an event occurring at time t in increasing an emotion of level e_t is represented by the linear operator $e_{t+1} = e_t + \Theta(1 - e_t)$. Here, Θ is a fraction between zero and one that determines how much the designated emotion changes toward its maximum value. One would expect that more important E-I rules would have larger associated values of Θ.

Let us assume that the activation coming from the variable in the E-I rule and its instantiated concept (i.e., from bully and Sam) add to a constant, say, one unit. Then the Θ value on the modulator valve tells us how much to increase the activation arriving at the parent emotion node (anger in this example) from the event in question. Thus, if the ongoing level of anger were .2 before this event, then the new anger level immediately after this event would be $.2 + .7(1 - .2) = .2 + .56 = .76$, which is a high level of anger. A built-in feature of an incremental emotion operator is that the model can accumulate a series of small provocations until it "boils over," as the parent anger node gains ascendancy over inhibiting emotions.[1]

Applications of the Model to Some Emotional Phenomena

We have now set forth the blackboard model, sketched how it arrives at an emotional appraisal, and illustrated the data structure it stores in memory about an emotional event. Let us now show how the model accounts for several emotional phenomena noted before.

State-Dependent Retrieval. Memory records emotional events along with instantiated E-I rules in data structures like that shown in Fig. 13.12. This system automatically indexes such data structures by the emotions felt at the time. Thus, if the person for any reason returns to the indicated emotional state (say, with

[1]The principle illustrated assigns "importance" to the E-I rule or participants in it. A complementary approach is for the system to have assigned importance weightings (in advance) to conceptual objects such as people ("Johnny"), property ("my car"), and aggregations ("my nation, church, race"). When activated, all concepts transmit a certain level of activation. One might suppose that this activation increases with the importance of the concept. On this account, the idea of one's self, spouse, or children would just automatically send out more activation once it was triggered, thus simulating greater "attentional demand." If so, and if this activation gets shunted via a modulator switch (see Fig. 13.12) into an emotion node, this means that more important concepts could lead to greater levels of emotional response when they participate in involving an E-I rule. Thus, in Fig. 13.12, if Johnny is especially beloved by the model, it should feel even greater sympathy. We have not yet thought through all the consequences of this "concept importance" approach to calculating emotional arousal. It seems to have some obvious good points as well as some flaws.

high-intensity activation of the parent anger node), the anger modulator valve in Fig. 13.12 will pass that activation into the memory record shown. Thus, the "Sam beat up Johnny" incident will be partially activated. The new insight here is that a given episode may be retrievable by cuing with either of the emotions (sympathy, anger or distress) felt toward the participants in the scene.

Priming Emotional Interpretations. Let us examine the consequences of activating a particular parent emotion node like anger. An obvious consequence is that this will send activation to all E-I rules that have anger as one of their modulator nodes. This amounts to "priming" those E-I rules into readiness for use in interpreting the next cognitive situation. Priming facilitates the system searching the E-I rules in order of their current level of activation, and selecting the first E-I rule that can be instantiated in the cognitive situation.

This priming causes the model to interpret ambiguous situations in a manner congruent with its prevailing emotion. Thus, the angry model has a "chip on its shoulder," is quick to take offense, is ready to find fault. Thus, if Sam were now to smile at the model who is angry at him, the model would interpret the smile as a mocking smirk and worthy of further anger.

An implication of this theory is that mood-congruent situations should be identified more quickly than mood-incongruent situations. For example, when judging the emotion displayed in a photograph of a face, the anger viewer should "see anger" quicker than sadness, whereas the sad viewer should "see sadness" quicker than anger.

Emotional Content of Fantasies. Our TAT results showed that people's feelings influenced the content of the imaginative stories they told. Although we are unable to develop a complete model of story construction, it is clear that a crucial early stage will be the storyteller's selection of a *theme*, such as achievement, or romance, or aggression and murder. Such high-level themes are linked to the parent emotion nodes. Thus, when the story constructor is feeling angry, themes reflecting frustration or aggression are primed and selected to control story construction.

Replaying an Emotional Memory. People can replay to some extent their original emotional reaction using the memory of a real or imagined event. How can this be explained? Our model says that to re-experience an emotion exactly, one must duplicate exactly the state of the blackboard at the time of that earlier experience. The more closely one can approximate the original state—the other ongoing feelings, the uncertainties of that moment—the more faithful will be the replay of that emotion. In replaying a memory, a propositional network structure like that in Fig. 13.12 is reactivated, placed on the blackboard, and the emotions designated in that memory record (by its modulator nodes) are "turned on."

The extent to which the parent emotion nodes are turned on again will depend on several factors that may have changed since the original experience. The most

conspicuous change is in people's uncertainty about what an event signified (e.g., how badly was Johnny hurt by the bully?); they were uncertain during the original experience, and that contributed to their arousal level then, but later we know the consequences of the event, so we cannot become so aroused by the event.

Reinterpreting While Replaying a Memory. When a memory is posted on the blackboard and replayed, that is the occasion for the model to automatically try to interpret the event once more. Although the event's memory record refers to an original E-I rule (P238 in Fig. 13.12), the system nonetheless searches for any E-I rules that will match the propositions contained in the memory record (i.e., that "Sam hit Johnny"). Under what circumstances will a reexamination of that event now come up with a different emotional interpretation? We consider several ways this can happen in the model:

1. Changed Mood. If the model is now feeling happy and loving, E-I rules refering to charity and forgiveness will be primed and will become available for quick matches to blackboard events. Thus, Sam's actions may be matched by an E-I rule that interprets Sam's hitting as part of a playful prank of exuberant children.

2. Changed Facts. Suppose the model has learned that Sam is a retarded person who has been picked on repeatedly by Johnny who is a nefarious scoundrel. Then the intermediate inferences of Fig. 13.12 are all wrong now. So a different E-I rule should now be instantiated, namely, one that causes one to feel satisfaction at retribution for injustice. That will cause the model to feel glad to see scoundrel Johnny get what he deserves.

3. Changed Importance. The model may have been upset originally at the incident because it liked Johnny and assigned importance to events affecting his welfare. Suppose that that liking changes into neutrality or even dislike. Then reevaluation of the hitting event in memory will lead to a greatly reduced emotional reaction. As a contrasting example, when Fred labeled his burglars as "murderers," his replaying of the burglary of his home would evoke greater discomfort.

The other component of importance was uncertainty: the initial hitting incident may have aroused more emotion because the model was uncertain how much harm was being done to Johnny, but this uncertainty is no longer present.

4. Changed E-I Rule. People occasionally change their E-I rules perhaps through some "conditioning" or "extinction" experiences or social influences (e.g., religious conversions). Thus, if the model joins the Nazi party and learns to admire strong superman and hate the weak, then new E-I rules will be learned. (Rule learning is modeled in ACT by different "strengths" of frequently reinforced versus neglected E-I rules.) In this case, when the model replays its memory of Sam hitting Johnny, that should now find a match to the LHS of a different E-I rule, which causes the model to feel admiration for Sam and contempt for Johnny. We have shown how in replaying a memory the model will search for an E-I rule with a matching LHS. When a match occurs, a token of that rule will be instantiated (if it is the same rule as before, this is a second instantiation of it). This invocation and instantiation of an E-I rule by a proposition on the blackboard is one of the causes of emotion in the model. If the E-I rule invoked during this reappraisal is different than originally, then different emotions will be evoked (partly at the same time as the old ones are replayed). Significantly, this new interpretation and its E-I rule will be stored in memory now with that old memory record. This gives the model the ability to update its feelings—to say of an event "That used to upset me, but it no longer does."

Conflicting Emotional Interpretations. Conflicting emotional interpretations may give rise to "mixed feelings" in many situations. For example, imagine that a beloved aunt dies and leaves you $5000; you are likely to feel grief over her death and happiness over your windfall. In our theory, this is represented simply as two events (the death and the inheritance) with two concurrent emotional interpretations. Because the interpretations give rise to different emotions, the intensity of each will be descreased by the other.

"Mixed feelings" are represented by simultaneously active and conflicting emotional interpretations. A related situation involves an active emotional interpretation that conflicts with remembered emotional interpretations. For example, if you hear that "Sam hit Johnny," but you like Sam very much, then you will moderate your anger toward him. There are several ways that this can happen. First, you may fail to infer that Sam is a bully, in which case the E-I rule that specifies anger toward bullies will not apply. This is quite plausible, especially given the powerful effect of mood on selective recall and inference. On the other hand, if you are already angry when you hear that Sam beat up Johnny, you may experience mixed feelings: anger toward Sam for his violence to Johnny, but also happiness from recalling one or more positive things about him.

Hot Versus Cold Cognition. It appears necessary that the long-term memory should have one unit or node that is the emotional feeling itself, another unit that is the concept of the emotion, and a third unit that is the name of that concept. See Fig. 13.13 for a diagram of this. The first node is responsible for the subjective emotional experience, the autonomic and expressive behaviors, and the cognitive influences; the second node collects together the meaningful connections of the concept (and points to the first node as its referent); and the third node (the "name") is used to talk about the concept and its referent emotion. The name of the concept can be activated and used without in the least turning on the referent feeling. This is done when we discourse dispassionately about an

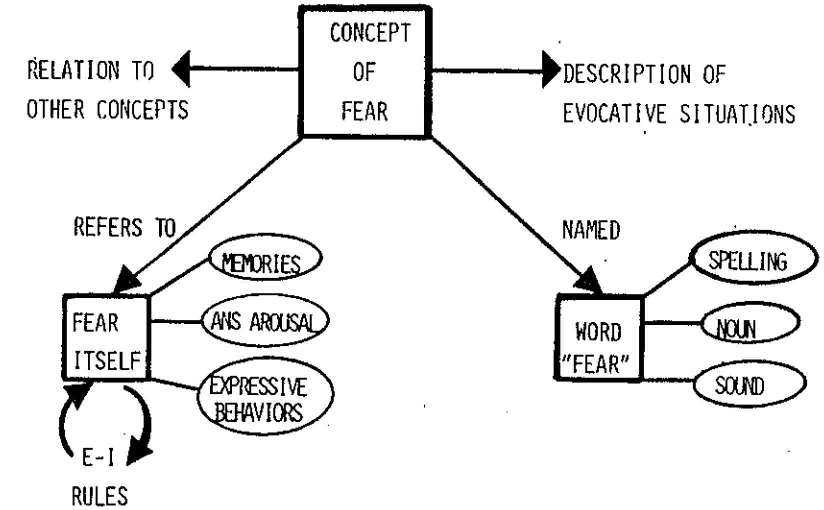

FIG. 13.13. Relations among network nodes representing an emotion, the concept of it, and words naming the concept.

emotion. The only natural way to "turn on" the feeling node is through instantiation of an E-I rule by a blackboard situation. (An unnatural method is to directly stimulate an emotion center with a brain electrode.); that is, natural emotions cannot be turned on directly but only through matching an E-I rule to an actual, imagined, or remembered situation represented on the blackboard. People cannot just feel sad or angry in the abstract; rather, they must activate a memory or imaginary scene that will lead them into sadness or anger.

This method for arousing emotions seems right intuitively. The technique is well-known among acting teachers (e.g., Stanislavski, 1936) and every fledgling actor or actress collects a set of emotional memories that they can press into service to stimulate the emotions their stage roles demand.

Reasoning about Emotional Memories. Recall that an emotional experience causes a data structure like that in Fig. 13.12 to be recorded in memory. This stores the event (Sam hit Johnny) along with the pattern of emotional reactions evoked and a token of the E-I rule whose instantiation led to that emotional reaction. In Fig. 13.12, the top-level node is a token of E-I rule P238 about bullies hurting weaklings.

This memory structure may be compared to that proposed in the earlier, simpler theory (see Fig. 13.4); that earlier theory would have just associated one emotion, say anger, with the whole episode by a direct linkage. The new approach provides for a more discriminating assignment of emotional reactions to participants in the episode. Perhaps more importantly, the new scheme permits the model to reason about or give justification for its emotional interpretations. We believe this is an important advantage that warrants further comment.

A first advantage of storing the instantiated E-I rule with the event's memory is that it enables one to reinterpret the event emotionally in case his E-I rules change. Thus, someone can replay an upsetting memory, notice that it is now best interpreted by an altered E-I rule, and thus store this new emotion with the event memory. The former memory model did not have the capability of "understanding" (giving reasons for) any such changes that the model would record.

Second, the "examined life" to which a culture exhorts its members requires that they be able to *justify* their acts, attitudes, and emotional reactions. Our new data structure can do this because it points to the E-I rule justifying its emotional reaction to the original event. Moreover, this particular E-I rule is itself located within a generalization hierarchy (imagine Fig. 13.11 generalized). Thus, a childhood bully hitting a weaker child is just a special case of a powerful person taking advantage of a weak one, which in turn is a special case of an act that exaggerates a preexisting inequity of happiness or utility (i.e., the strong get stronger). Thus, if asked repeatedly to justify its wrath at bully Sam, the model can be "backed up" its belief hierarchy, where it will provide ever more general rationales for its indignation. In moral arguments, people are often called upon to justify their moral ("emotional") judgments in just such abstract terms.

Third, this generalization hierarchy of emotional rules gives one a way to conceptualize how cognitive therapy affects emotional disorders. For example, in rational emotive therapy as practiced by Albert Ellis (1962), the maladjusted patient is conceived to have certain irrational beliefs, expectations about himself, other people, and their relationship to him. Sample irrational beliefs can be represented as very general rules such as:

IF I'm not perfect,
THEN I am worthless and should be depressed.

IF everyone doesn't love everything I do,
THEN I am worthless and should be depressed.

IF everything doesn't go well,
THEN I'm to blame and should feel guilty.

In rational emotive therapy and in Beck's (1976) cognitive therapy, much therapeutic effort is concentrated on unearthing the patient's irrational beliefs, then challenging and changing them, replacing them with less-self-defeating beliefs. The basic premise of the therapy is that emotional appraisals can be viewed as the application of abstract rules to specific situations, and that it is possible to change these rules directly and thus effect therapeutic changes in maladaptive emotional syndromes. The therapy is undeniably effective in many instances.

Our theory of an E-I rule hierarchy provides a way to discuss the effect of such "attitude" therapy. The problem encountered in such therapy is also predicted by the theory, namely, that a general rule freshly taught by the therapist will be short-circuited (not applied) by the client reverting to specific, low-level E-I rules of long, strong standing. Much guided practice *in vivo* in using the new emotional interpretation rules is needed to bring about therapeutic changes.

Affective Similes and Metaphors. The hierarchical organization of E-I rules permits one to compare episodes and note their similarity in terms of the general E-I rule and pattern of affect they evoke. Thus, a man beating his dog is like Sam hitting Johnny, and both are like someone pulling the wings off a butterfly. All are instances of E-I rules that identify various forms of cruelty, and the specific E-I rules are collected under a common, generalized cruelty E-I rule. Because of their common affective response and invocation of common E-I rules, one of these episodes may remind us of another. Thus, the Wehrmacht army's blitzkreig into Belgium in 1940 may remind us of a bully's overpowering a defenseless opponent. These sorts of remindings by affective similes occur frequently in literary metaphors and in everyday conversations.

To summarize this section, by storing instantiations of reasoning rules with the episodes in which they are used, the model can reason intelligently at a later time about the episode. It can explain how it felt, and can justify its feelings at several levels of abstraction. It can compare one emotional episode to another, even if they do not share concepts, on the basis of common or related E-I rules. Importantly, it can also say why it feels differently about an episode at a later time; for example, if it learns that the man was not actually beating the dog but was only pretending to beat it for a circus routine, the model can reinterpret the episode and by comparing the original E-I rules with the more recent ones, it can say why it feels differently.

These conditions do not detract in any way from the advantages of the original semantic network model. Those network ideas are still used to account for state-dependent memory and selective learning by the usual mechanisms of cuing and spreading activation. The new model adds capabilities for appraising emotional situations and reasoning about those appraisals.

Final Comments

Let us summarize briefly what we think we have accomplished. First, we reviewed experimental results showing diverse influences of emotion upon selective learning, selective retrieval, and emotionally biased judgments and interpretations. The effects are strong, wide-spread, and intuitively understandable. Second, we reviewed the earlier semantic network theory of emotional memories that sufficed to explain roughly most of those experimental results. Third, we noted incompletenesses in that theory and pointed to the need for an emotional appraisal system. Fourth, we introduced the blackboard control structure for coordinating the knowledge sources we believe are involved in emotionally appraising situations. We described some of the knowledge sources and sketched briefly how they might interact on the blackboard in arriving at an emotional interpretation. Fifth, we then applied this blackboard model to several complex phenomena in emotion such as replaying and reinterpreting one's earlier emotional experiences.

We would emphasize that these latter ideas are in an early stage of development, and we would feel more secure if we had an operating computer program that implements our verbal ideas. We are impressed by the difficulty of the problems to be faced by any theory of emotional behavior. We hope that we have contributed in some measure to the clarification or resolution of some of the problems.

ACKNOWLEDGMENT

This research was supported by Grant MH–13950 from the National Institutes of Mental Health to G. H. Bower.

REFERENCES

Anderson, J. R. *Language, memory and thought*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1976.

Bartlett, J. C., Burleson, G., & Santrock, J. W. Emotional mood and memory in young children. *Journal of Experimental Child Psychology*, in press.

Beck, A. T. *Cognitive therapy and the emotional disorders*. New York: International Universities Press.

Bower, G. H. Mood and memory. *American Psychologist*, 1981, *36* (No. 2), 129–148.

Bower, G. H., Gilligan, S. G., & Monteiro, K. P. Selectivity of learning caused by affective states. *Journal of Experimental Psychology: General*, 1981, *110*, 451–473.

Bower, G. H., Monteiro, K. P., & Gilligan, S. G. Emotional mood as a context for learning and recall. *Journal of Verbal Learning and Verbal Behavior*, 1978, *17*, 573–585.

Carbonell, J. G. *Subjective understanding: Computer models of belief systems*. New Haven, Conn.: (A.I. Tech. Rep. 150), Yale Computer Science Dept. 1978.

Clark, M. S., & Isen, A. M. Toward understanding the relationship between feeling states and social behavior. In A. Hastorf & A. M. Isen (Eds.), *Cognitive social psychology*. Amsterdam: Elsevier, 1982.

Eich, J. E. The cue-dependent nature of state-dependent retrieval. *Memory and Cognition*, 1980, *8*, 157–173.

Eich, J. E., Weingartner, H., Stillman, R. C., & Gillin, J. C. State-dependent accessibility of retrieval cues in the retention of a categorized list. *Journal of Verbal Learning and Verbal Behavior*, 1975, *14*, 408–417.

Ellis, A. *Reason and emotion in psychotherapy*. New York: Lyle Stuart, 1962.

Gilligan, S. G., & Bower, G. H. *Emotional mood and remembering one's autobiography*. First draft of manuscript, Department of Psychology, Stanford University, 1981.

Henry, G., Weingartner, H., & Murphy, D. L. Influence of affective states and psycho-active drugs on verbal learning and memory. *American Journal of Psychiatry*, 1973, *130*, 966–971.

Isen, A. M., Shalker, T. E., Clark, M. & Karp, L. Affect, accessibility of material in memory, and behavior: A cognitive loop? *Journal of Personality and Social Psychology*, 1978, *36*, 1–12.

Mowrer, O. H. *Learning theory and behavior*. New York: Wiley, 1960.

Plutchik, R. *A psychoevolutionary synthesis*. New York: Harper & Row, 1980. (a)

Plutchik, R. A language for the emotions. *Psychology Today*, February 1980, *2*, 68–78. (b)

Roseman, I. *Cognitive aspects of emotion and emotional behavior*. Paper presented at the meeting of the American Psychological Association. New York: September 1979. (Unpublished manuscript, Department of Psychology, Yale University.)

Schiffenbauer, A. I. *The effects of a judge's affective state upon his judgment of the emotional state of others*. Unpublished doctoral dissertation, Department of Psychology, Stanford University, 1970.

Tversky, A., & Kahneman, D. Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 1973, *5*, 207–232.