Dynamic Memories:
Analysis of an Integrated Comprehension and Episodic Memory Retrieval Model

Trent E. Lange
Artificial Intelligence Lab, Computer Science Dept.
University of California, Los Angeles
Los Angeles, CA 90024

Charles M. Wharton
Department of Psychology
University of California, Los Angeles
Los Angeles, CA 90024

Abstract

Most AI simulations have modeled memory retrieval separately from comprehension, even though both activities seem to use many of the same processes. We have developed REMIND, a model that performs both episodic memory retrieval and language understanding with a single spreading-activation mechanism. This approach has a number of advantages over retrieval-only models. First, because the comprehension process makes inferences about actors* plans and goals, REMIND is able to get abstract remindings that would not be possible without an integrated model. It also allows a more psychologically-plausible model of reminding than previous approaches, since all aspects of a text's interpretation affect what is retrieved through the spreading-activation process, as in human reminding. An inferencing-based retrieval model such as REMIND also has several computational advantages over pure retrieval models. The effects of the understanding process eliminate the need for the separate, purely structural comparisons used in most analogical retrieval models. Further, it potentially explains how the explicit indexing of case-based reasoning models can be eliminated, while retaining its benefits as an emergent property of the comprehension process.

1 Introduction

The most parsimonious account of comprehension and reminding is that they "amount to different views of the same mechanism" [Schank, 1982]. However, most AI models that perform memory retrieval do so in isolation from the language understanding process. Different retrieval models approach varying aspects of the retrieval problem and have different goals, but nearly all are given fully hand-coded representations of the memory episodes (or cases) they use.

Most psychological models of memory retrieval simulate empirical results showing that reminding is based on both surface feature similarities (e.g., shared words) and analogical similarities (e.g., shared inferences and themes) [Wharton et al, in press]. Systems such as ARCS [Thagard et a/., 1990] and MAC/FAC [Gentner & Forbus, 1991] model this with two-stage retrieval processes that first search for all episodes sharing surface features with the cue. They then select the episode that shares the most surface and analogical similarities by mechanisms that explicitly calculate structural isomorphism (or analogical similarity) between the cue and targets.

Case-based reasoning (CBR) models (cf. [Riesbeck & Schank, 1989; Owens, 1989]) recognize pragmatically-useful index patterns that allow retrieval of episodes (or cases) likely to aid their current task. CBR models generally simulate reasoning of experts within a given domain, rather than general human reminding. As in general reminding models, most CBR models use hand-coded cases and operate separately from the comprehension process, though a few do some rule-based reasoning to make explanations (cf. [Hammond, 1989]). Most CBR models instead focus on deciding which particular abstract indices (or features) are most useful for retrieval of useful cases in different problem-solving tasks.

Although separating memory retrieval from the language understanding process makes accounts of the phenomena more manageable, it is undeniable that real-world retrieval results from comprehension processes. We have developed REMIND (Retrieval from Episodic Memory through INferencing and Disambiguation) [Lange & Wharton, in press], a model that integrates language understanding and memory retrieval in a single structured spreading-activation mechanism. This approach has several computational advantages over retrieval-only models and provides a more psychologically-plausible model of reminding. This paper shows several simulations that illustrate some of these advantages, contrasting it with retrieval-only and CBR approaches.

2 Overview of REMIND

REMIND is initially given a syntactic representation of a short text as a cue. Using general knowledge stored in its long-term memory, REMIND constructs an elaborated interpretation of the cue and retrieves the episode that is most similar to the surface and inferred features of that representation.

REMIND's structured spreading-activation networks encode world knowledge about concepts and general knowledge rules for inferencing in the same way as ROBIN [Lange & Dyer, 1989], a structured connectionist model that performs high-level inferencing and disambiguation for natural language understanding. Structured connectionist models seem to be particularly well-suited to language understanding because their constraint satisfaction abilities are ideal for integrating contextual evidence to perform disambiguation and priming, while their network structure allows for the representation of complex knowledge and parallel inferencing.

REMIND's networks also contain representations of prior episodes, such as Fred put his car in the car wash before his date with Wilma (Car Wash) and Billy put his Playboy under the bed so his mother wouldn 7 see it and spank him (Dirty Magazine). The representations used are the actual plan/goal analysis (or interpretation) that was inferred by REMIND when input for them was first presented to the network. These prior episodes are indexed into the semantic comprehension network through connections with all the knowledge structures with which they were understood. To perform retrieval, REMIND is given a short text passage to understand and use as a cue. This understanding process often requires disambiguation and for a number of inferences to be made. Consider:

John put the pot inside the dishwasher because the police were coming. (Hiding Pot).

Although it initially appears that John is cleaning a cooking pot, this sentence is disambiguated and interpreted to mean John was hiding marijuana from the police to avoid being arrested. To understand such cues, units in the network representing the cue and its syntactic bindings are clamped to high levels of activation. Activation is then spread through the network. By propagating signature activation patterns [Lange & Dyer, 1989], the network makes the different possible inferences explaining the input in a manner similar to marker-passing systems (cf. [Riesbeck & Martin, 1986]). For example, one of the multiple interpretation paths that gets inferred (and activated) as a possible explanation for Hiding Pot is the interpretation that John was trying to hide the pot from the police to satisfy his goal of avoiding arrest. Other interpretations concurrently activated include the possibilities that he was trying to clean the pot or store it. Activation spreads until the network settles.The units with the most activation represent the most plausible set of inferences and the network's disambiguated plan/goal interpretation of the cue.

Because units representing long-term memory episodes are connected within the network, episodes having concepts related to the elaborated cue also become highly activated. This includes episodes only superficially-related to the cue because of surface feature overlap (e.g., episodes involving police or illegal drugs) and episodes related abstractly because they share similar inferred plans and goals (e.g. episodes sharing the inferences that a person was trying to Avoid-Detection of something to avoid a Punishment). After the network settles, the episode receiving the most activation from the cue's interpretation becomes the most highly activated, and is retrieved as the best match for the cue.

Thus, in REMIND, a single mechanism drives both language understanding and memory retrieval processes. The same spreading-activation mechanism that infers a single coherent interpretation of a cue also activates episodes retrieved from memory. Episodic activation results from both the surface semantics of the input (i.e., different possible word and phrase meanings) and the deeper thematic inferences made from the input. Accordingly, recalled episodes depend on both surface and analogical similarities with the cue. Because both inferencing and memory retrieval occur within a single integrated network, the context in which interpretations are formed affects the episodes that are retrieved, which in turn influences the context in which disambiguation and interpretation of input takes place. Thus, text comprehension and memory retrieval processes are tightly coupled.

3 Cue Understanding in REMIND

As with ROBIN, REMIND uses structured networks of simple connectionist units to encode semantic networks of frames and rules representing world knowledge, such as the scripts, plans, and goals [Schank, 1982] necessary for understanding stories in a limited domain. Figure 1 shows a partial overview of a REMIND network and the knowledge given to it (by hand, as in most structured models). However, it is given no specific information about episodes it will understand.

Knowledge given to REMIND is used to construct the actual structure of the network before any processing begins. As with other structured connectionist models, nodes in the network represent particular frames or roles. Relations between concepts are represented by weighted connections between nodes. Activation on concept nodes is evidential, corresponding to the amount of evidence available in the current context. The network also has additional structure to solve connectionist networks' variable binding problem by propagating signature activation patterns representing bound concepts [Lange & Dyer, 1989]. The network makes inferences in parallel by propagating signature bindings (such as of John, Marijuana, and Cooking-Pot) over connections between binding units that represent general knowledge rules.

As an example, consider how Hiding Pot is understood by the network. To represent John put the pot inside the dishwasher, Transfer-Inside is clamped to a high level of evidential activation (black box in Figure 1). The binding units of its roles (not shown) are clamped to the signature activations of its bindings (John for its Actor, Cooking-Pot or Marijuana for its Object, and Dishwasher for its Location). These signature bindings then propagate, as activation, over connections to the corresponding roles of neighboring frames. This propagation allows the network to infer that the pot is Inside-Of the dishwasher and that it was done either because it was going to be cleaned (Inside-Of-Dishwasher and following frames) or because it would be blocked from sight (Inside-Of-Opaque and following frames).

As signatures propagate to perform inferencing, evidential activation spreads and accumulates along conceptual nodes to disambiguate between competing inferences. Initially the Inside-Of-Dishwasher path receives the most activation because of feedback between its strong stereotypical connections to Cooking-Pot and Dishwasher. However, activation feedback between Inside-Of-Opaque and inferences from the police coming (Transfer-Self...Block-See) and the Police-Capture frames causes Inside-Of-Opaque to end up with more activation than Inside-Of-Dishwasher and Marijuana with more activation than Cooking-Pot.

The network's final interpretation of Hiding Pot includes the most highly-activated path of frames in Figure 1 and their network signature bindings. This interpretation includes the inferences that (a) Marijuana is inside of an opaque dishwasher (Inside-Of-Opaque) and is blocked from sight (Block-See), (b) John possesses illegal marijuana (Possess-Illegal-Obj), and (c) John is in danger of being arrested by the police (Police-Arrest). Note that alternative interpretation paths retain activation for possible reinterpretation, since REMIND uses inhibition that normalizes activations rather than driving losers to zero. See Lange & Dyer [1989] for further details on how the network performs such inferencing and disambiguation for Hiding Pot and other inputs.

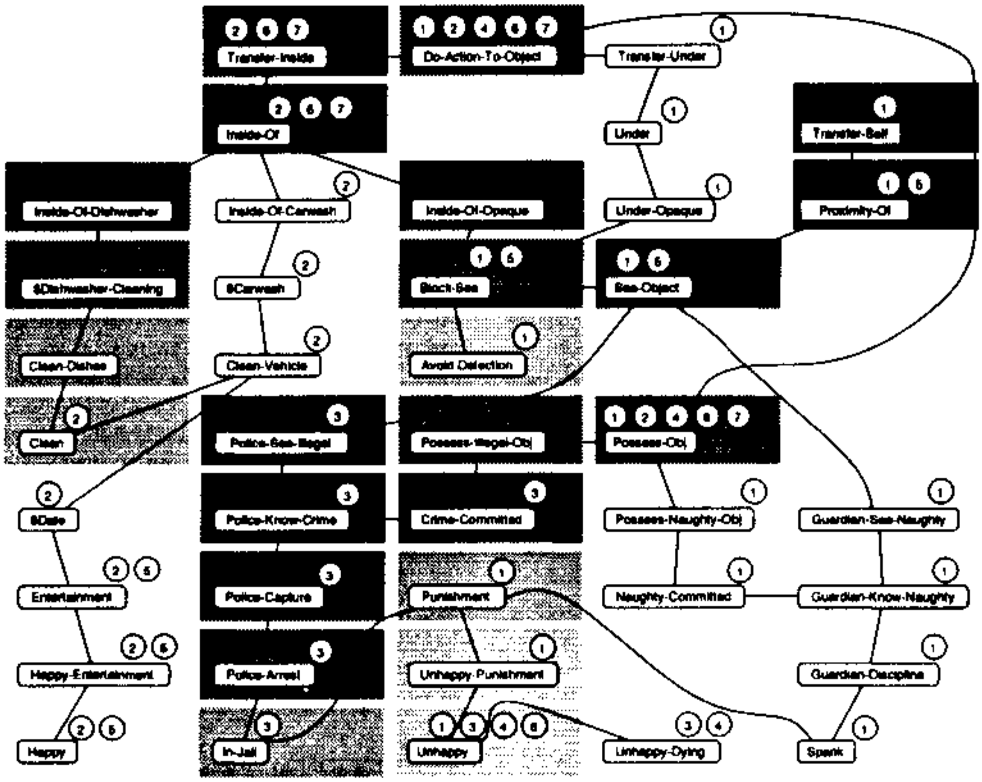

Figure 1. Overview of network segment after activation has settled in processing Hiding Pot. Actual network represents frames and their roles by structured sub-networks of units holding evidential and signature activations. Gray boxes represent level of evidential activation on the frames (darker = higher activation). Circles above frames indicate a long-term instance of frame in an episode. Episodes understood and stored here: 1) Dirty Magazine. 2) Car Wash. 3) *Jane shot Mark with a Colt-45. He died.* 4) *Betty wanted to smoke a cigarette, so she put it on top of the stove and lit it.* 5) *The pleasure boat followed the whales to watch them.* 6) *Barney put the flower in the pot, and then watered it.* 7) *Mike was hungry. He ate some fish.* 8) *Suzie loved George, but he died. Then Bill proposed to her. She became sad.*

4 Memory Retrieval

In REMIND, memory retrieval is a side-effect of the spreading-activation understanding process. Representations of previously-understood episodes are connected directly to the semantic network that understood them originally. Knowledge structures activated when understanding a cue activate similar episodes that were stored in the network earlier. This direct form of "indexing" causes episodes that share many conceptual similarities with the cue to become active during interpretation. The most active episode is retrieved.

4.1 Network Encoding of Episodes

Whereas REMIND uses hand-coded general world knowledge, it is not given any information about the particular episodes it processes and stores in long-term memory. Target episode representations are created entirely by REMINDS spreading-activation understanding process. Input for each episode is presented to the network, which infers an interpretation by the spread of signature and evidential activation as described before. Next, units and connections are added to store the episode* s entire resulting interpretation in the network. Thus, each episode's representation includes all aspects of its interpretation, from its disambiguated surface features (such as the actors and objects in the story) to the plans and goals the network inferred the actors were using.

For example, consider how Dirty Magazine *(Billy put the Playboy under his bed so his mother wouldn 't see it and spank him)* is processed and encoded as a memory episode. First, input for its phrases is clamped and an interpretation inferred. As in Hiding Pot, the network infers that somebody is hiding something (Avoid-Detection) and that it is blocked from sight (Block-See). Here, however, the inferred signatures show that it is Billy hiding a Playboy-Magazine rather than John hiding Marijuana. Several other knowledge structures involved in Hiding Pot (e.g. Proximity-Of, Possess-Obj, Punishment) are also activated by Dirty Magazine. However, there are a number of differences, e.g. frames of the Guardian-Discipline structure are part of Dirty Magazine's interpretation, but the Police-Capture frames are not.

Once an interpretation is formed of an episode, units and connections are added to the network (by hand) to represent all of its instantiated frames and elements. [Lange & Wharton, in press] describes the actual units and connections added to do this; the important part is that the added units representing each instantiated frame (such as a particular instance of Avoid-Detection where Billy was hiding a Playboy-Magazine) are added *locally* to its semantic network frame and that their connections encode their bindings. The new instance units are also all connected to an episode unit that groups all of the episode's elements. An overview of the final result is shown in Figure 1, with the nodes having a circled " 1 " above them representing the frames inferred and encoded as part of Dirty Magazine's representation. Other circled numbers represent elements of other stored episodes' interpretations.

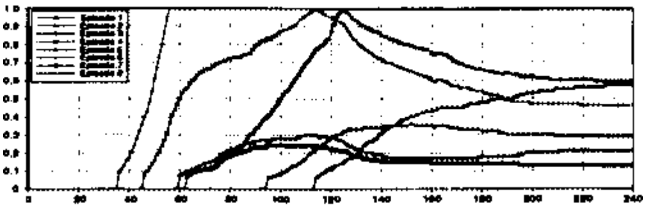

Figure 2. Episode unit activations after presenting **Hiding Pot.**

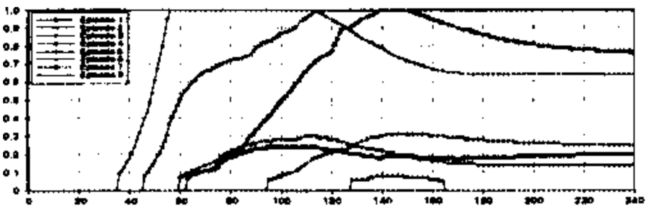

Figure 3. Episode unit activations after presenting **Dinner Party**.

4.2 The Retrieval Process

With episodes understood and stored in the network, retrieval is performed simply by presenting an input cue to the network to be understood. Because instance units representing episodes are connected directly to the normal semantic units, they become activated by the inferencing spread of activation. The more similarities an episode shares with the inferred interpretation of a cue, the more of its instances become active and the more activation its episode unit receives.

Figure 2 shows activations of the eight episodes from Figure 1 during understanding of Hiding Pot. Episode.6 (Barney put the flower in the pot, and then watered it) initially becomes highly active because it shares a number of surface features — e.g. both involve a Transfer-Inside, and Planting-Pot gets activation from pot. Similarly, Episode.2 and episodes having varying degrees of shared features become active. However, as time goes on, the hiding and punishment frames are inferred and become active. Episode.1 (Dirty Magazine)'s activation thus climbs and eventually wins, because it shares the most surface and abstract features of any episode with Hiding Pot's interpretation (see Figure 1). It is therefore retrieved as the episode most similar to Hiding Pot.

5 Experiments

REMIND has been tested with a knowledge base having 206 conceptual frames and 333 inference rules. It has understood and retrieved the examples here and a number of other episodes of similar complexity. Here we briefly describe three additional simulations that illustrate (1) the importance of inferences and disambiguation on retrieval, (2) the strong influence of superficial feature similarities on retrieval, and (3) the effect of episodic recall on the understanding process.

5.1 Importance or Inferencing

An example of how strongly the inferencing and disambiguation of the model affects retrieval is shown in Figure 3, which shows activations after presentation of input for John put the pot inside the dishwasher because company was coming (Dinner Party). Note that although this cue differs from Hiding Pot by only a single word (company instead of police), the interpretation REMIND reaches is completely different (i.e. that he was cleaning a cooking pot to prepare for a dinner party). This causes a different episode to be recalled, Episode.2 (Car Wash), which shares the goals of cleaning something to prepare for an entertainment event.

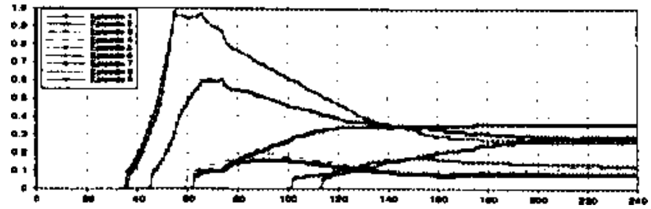

Figure 4. Episode unit activations of Figure 1 episodes and Episode.9 after presenting Hiding Pot.

5.2 Superficial Similarities

As in human reminding, REMIND often retrieves superficially similar episodes to a cue even when a better analogy exists. As an example, Figure 4, shows the episode activations after Hiding Pot is presented to a network having the eight episodes of Figure 1 and an additional superficially similar episode, Cheech put the grass inside the hong because Chong was coming (Episode.9). Notice that the activation of this new episode quickly dominates the others because of the surface features it shares with Hiding Pot. This superficially similar, but thematically dissimilar episode is therefore retrieved even though Dirty Magazine is a better analogy.

5.3 Effect of Reminding on Interpretation

REMIND's integration of the reminding and understanding processes has pragmatically interesting and useful effects on the understanding process. Episodes that become active while inferencing feed activation back into the comprehension part of the network. This can prime and bias the interpretation REMIND settles on for a given input. Consider:

> The star loved the plumber, but he was shot by a thief. Then the astronomer proposed to her. She started to cry. (Astronomer Proposal).

Two possible reasons for the movie star starting to cry are that the proposal either made her extremely happy (Happy-Proposal) or extremely sad (Unhappy-Proposal). Perhaps the most likely reason for her crying was that the proposal reminded her of murdered lover, therefore making her sad. REMIND, however, does not have the complex knowledge about memories and how they affect people's emotions thai would be necessary to make that interpretation. As shown in Figure 5, REMIND therefore ends up interpreting the star's crying in Astronomer Proposal as a Happy-Proposal, because of its weights' strong biases that marriages are happy. As Astronomer Proposal illustrates, REMIND often arrives at counter-intuitive interpretations of stories when the biases of its connection weights are too strong or when it does not have enough knowledge to make the needed inferences for the right interpretation. However, when there is a highly-analogous episode (or case) in memory, the influence of episodic retrieval upon text understanding can lead REMIND to a correct interpretation of its input. For example, consider:

> Suzie loved George, but he died. Then Bill proposed to her. She became sad. (Sad Proposal)

Sad Proposal is similar to Astronomer Proposal, but explicitly states that Suzie became Unhappy after the proposal.

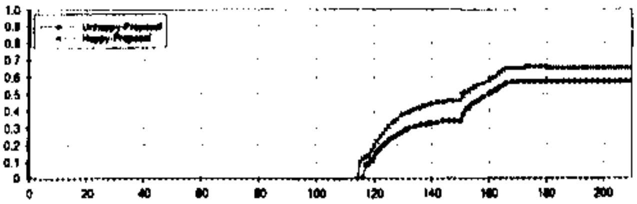

Figure 5. Activations of Happy-Proposal and Unhappy-Proposal interpretations after Astronomer Proposal is presented.

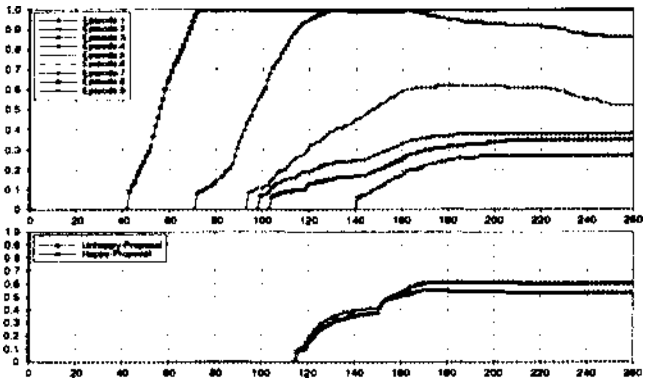

Figure 6. Activations of episodes and of Happy-Proposal and Unhappy-Proposal interpretations after Astronomer Proposal is presented to network containing Sad Proposal episode.

This leads the network to make the correct interpretation, that the Marriage-Proposal after the death of her lover was an Unhappy-Proposal. This interpretation, including the inference Unhappy-Proposal.8, is stored in memory as Episode.8 in Figure 1. Figure 6 shows the activations of Sad Proposal (Episode.8) and the other episodes as Astronomer Proposal is then understood by REMIND. As expected, Sad Proposal quickly dominates the other episodes because it is so similar to Astronomer Proposal Episode.3 becomes temporarily active because it also involves a shooting. However, Sad Proposal eventually wins and is retrieved.

The most interesting result in Figure 6 is the activation levels of the competing Happy-Proposal and Unhappy-Proposal frames. As when Astronomer Proposal was presented to the network without any episodes in memory, Happy-Proposal initially has more activation than Unhappy-Proposal. In this case, however, Episode.8 is highly active, and with it Unhappy-Proposal.8. Since active episode instances feed activation back into their concepts in the understanding network, Unhappy-Proposal gets significant activation from Unhappy-Proposal.8. This added evidence allows its activation to climb over Happy-Proposal, which gets no added evidence from memory episodes. Unhappy-Proposal therefore ends up with more activation than Happy-Proposal, so REMIND's interpretation is that the proposal made the movie star unhappy. The network therefore selects the correct interpretation of Astronomer Proposal because of activation feedback from an analogous case in memory, Sad Proposal.

REMIND's use of one spreading-activation mechanism for both comprehension and memory retrieval shows how recall can subtly affect interpretation. When stored episodes are similar to a cue that REMIND is interpreting, they feed activation back into the inferencing network. This feedback can bias REMIND's interpretation to be consistent with the active episodes, a limited form of case-based reasoning.

6 Comparison to General Reminding Models

REMIND is most directly comparable to ARCS [Thagard *et al,* 1990] and MAC/FAC [Gentner & Forbus, 1991], two other models of general, non-expert reminding. All three take into account psychological evidence showing that memory retrieval is strongly influenced by both surface and thematic similarities between a cue and episodes in memory. In contrast to REMIND, ARCS and MAC/FAC do not model the language understanding process, concentrating solely on retrieval of hand-coded targets. This allows them to retrieve more complicated episodes than REMIND can currently understand. On the other hand, REMIND's inferencing-based theory of retrieval has two significant advantages over retrieval-only models: (I) it shows how understanding and reminding can be modelled with a single spreading-activation mechanism, and (2) that this integration eliminates the need for the separate structural comparison mechanisms used to by ARCS and MAC/FAC to allow analogical retrieval.

A major criticism of ARCS and MAC/FAC is that neither model specifies how the representation of its input cues and episodes is formed or what kinds of knowledge those representations should generally include. Should the cue representations include only the surface propositions directly stated in a cue's text? Or should they include a fully elaborated interpretation of the cue, including a complete causal plan/goal analysis of the text and any abstract themes it involves, as in REMIND? Memory retrieval often cannot be performed without such inferences, as illustrated by examples in this paper and in CBR models. However, even if retrieval-only models such as ARCS and MAC/FAC were given fully elaborated cues, we believe that not modeling *the process* by which these representations are formed misses important features of reminding. Understanding of text varies depending on whether it is simply being skimmed or is being read carefully for its deeper ramifications. Thus, when the comprehension process is not modeled, there is no way to simulate the specific circumstances under which understanders infer and can use planning or thematic information in probing memory (cf. [Seifert *et* a/., 1986]). In contrast, REMIND explicitly models the cue interpretation process, and so can potentially explain when elaborated inferences are available to affect reminding.

REMIND is fundamentally different from ARCS and MAC/FAC in how it models the influence of analogical similarity on retrieval. Both ARCS and MAC/FAC model analogical reminding by explicitly factoring in the degree of *structural isomorphism* (or *relational consistency)* between the cue and targets into their best match computation. Isomorphism can best be explained by an example from [Thagard *et al.,* 1990] for the cue *The dog bit the boy and the boy ran away from the dog* (Boy Run). Compare this to the analogs *Fido bit John and John ran away from Fido* (John Run) and *Rover bit Fred and Rover ran away from Fred* (Rover Run). John Run is structurally isomorphic with Boy Run, because mapped objects play the same roles in mapped predicates. In both cases, the dog did the biting and the person it bit did the running. In Rover Run, however, it was the dog that ran from the person it bit. John Run is more isomorphic to Boy Run than is Rover Run, and is therefore a better analog.

Analogical similarity is hypothesized by ARCS and MAC/FAC to exert its effect on memory retrieval as a direct result of a specifically computed degree of syntactic isomorphism between cues and memory episodes. REMIND, in contrast, never explicitly computes the degree of isomorphism. Instead, relationally consistent targets are retrieved over relationally inconsistent targets only indirectly, when the different syntactic structure of each input *leads to different inferences.* For example, if presented with John Run, REMIND would infer that the boy ran away because he was afraid that dog would continue its attack. However, if presented with Rover Run, REMIND would infer that the dog ran away because it feared retaliation from the boy. Because of these different interpretations of the two episodes, REMIND would also retrieve John Run when presented with Boy Run as a cue. Unlike ARCS and MAC/FAC, however, REMIND does so without having to go through a separate stage to explicitly compute the degree of syntactic isomorphism.

We believe that the effects of syntactic isomorphism and relational consistency on memory retrieval can be fully explained by the understanding process. Relationally consistent episodes tend to have more similar inferences, interpretations, and themes than relationally inconsistent episodes. In REMIND, explicit measures of syntactic isomorphism are unnecessary, since analogical reminding occurs as a natural side-effect of interpreting and disambiguating an input text.

7 CBR Models and Indexing

Most case-based reasoning models are meant to be models of expert reminding and problem-solving in given domains. Unlike REMIND, they are not meant to be models of general, non-expert human reminding. An advantage of case-based reasoning models over REMIND is that their use of symbolic processing abilities allows them to handle longer and more complex episodes than REMIND (and connectionist models in general) can currently handle. On the other hand, as a model of comprehension and general reminding, REMIND is better able to explain psychological results such as the relatively high prevalence of remindings based on superficial similarities and on how the reminding and language understanding processes interact and effect each other.

There are many similarities between retrieval of episodes in CBR models and in REMIND. One of the major goals of CBR researchers is to find the indices that will enable retrieval of the cases most likely to help their current task (the *indexing problem).* Pragmatically useful indices usually include features such as abstract plans, goals, themes, explanations, and anomalous situations, depending upon the problem being solved. Although REMIND does not currently recognize all these types of indices, all are features that an ideal model would have to recognize and that would therefore be used for comprehension and retrieval in future versions of REMIND.

One of the advantages of REMIND's approach to storing episodes is that it avoids CBR models' indexing problem. Because episodes are simply stored (indexed) under *all* of the features that played a part in understanding them, there is no need for a separate computation stage to determine precisely *which* features they should be indexed under for best retrieval. The network's massively-parallel comprehension process eliminates the need to limit the number of indices used in order to constrain search time. Further, this approach has the advantage of allowing episodes to be retrieved in contexts and situations other than those the problem-solver (or index evaluator) originally considered useful.

Many of the desirable features of CBR indexing methods emerge from the dynamics of the spreading-activation process and how episodes are learned over time. For example, one important feature of a useful index is how unique it is. Although REMIND indexes its episodes under all of their features, relatively unique features affect retrieval more than common ones simply because they activate fewer episodes (compare Possess-Obj to the more abstract Avoid-Detection and Punishment frames in Figure 1). Another important aspect of the spreading-activation process is that particularly salient features receive the most activation and therefore automatically act as stronger retrieval indices. While it does not currently approach the problem-solving of many CBR models, an extension of REMIND that did so would also focus activation on the types of problems and failures being examined, therefore naturally emphasizing useful indices. Such a model might show how the benefits of explicit indexing in CBR models can fall out of the comprehension process.

8 Conclusions

REMIND's use of a single spreading-activation mechanism to perform both comprehension and retrieval ensures that the features inferred from a cue during understanding will access episodes in memory that share similar inferences. This integration of comprehension and retrieval is a more psychologically-plausible way of producing analogical reminding than previous models. It also has several computational advantages. Because the inferencing process activates abstract plans and themes, the explicit structural isomorphism computations needed to allow retrieval of analogies in retrieval-only models such ARCS and MAC/FAC are unnecessary. REMIND's massively-parallel approach to comprehension and encoding also potentially explains how the explicit indexing of CBR models can be eliminated, while retaining its benefits as an emergent property of the comprehension process

References

Centner, D., & Forbus, K. (1991). MAC/FAC: A model of similarity-based retrieval. *Proceedings of 13th Conf. of Cognitive Science.*

Hammond, K. (1989). *Case-Based Planning.* Academic Press.

Lange, T., and Dyer, M. (1989). High-level inferencing in a connectionist network. *Connection Science,* 1:181-217.

Lange, T., &. Wharton, C. (in press). REMIND: Retrieval from episodic memory by inferencing and disambiguation. In Barnden and Holyoak (eds.), *Advances in Connectionist and Neural Computation Theory, Vol. 2: Analogical Connections.* Norwood, NJ: Ablex.

Owens, C. (1989). Integrating feature extraction and memory search. *Proceedings of the 11th Conf of Cognitive Science Society.*

Riesbeck, C. and Martin, C. (1986). Direct memory access parsing. In J. Kolodner and C. Riesbeck (eds). *Experience, Memory, and Reasoning,* 209-226. Lawrence Erlbaum.

Riesbeck, C, & Schank, R. (1989). *Inside Case-Based Reasoning.* Hillsdale, NJ: Lawrence Erlbaum.

Schank. R. (1982). *Dynamic Memory.* New York: Cambridge Press.

Seifert, C, McKoon, G., Abelson, R. & Ratcliff, R. (1986). Memory connections between thematically similar episodes. *Journal of Exp. Psychology: Human Learning and Memory,* 12, 220-231.

Thagard, P., Holyoak, K., Nelson, G., & Gochfeld, D. (1990). Analog retrieval by constraint satisfaction. *Artificial Intelligence* 46:259-310.

Wharton, C, Holyoak, K, Downing, P., Lange, T., Wickens, T. & Mclz, E. (in press). Below the surface: Analogical similarity and retrieval competition in reminding. *Cognitive Psychology*