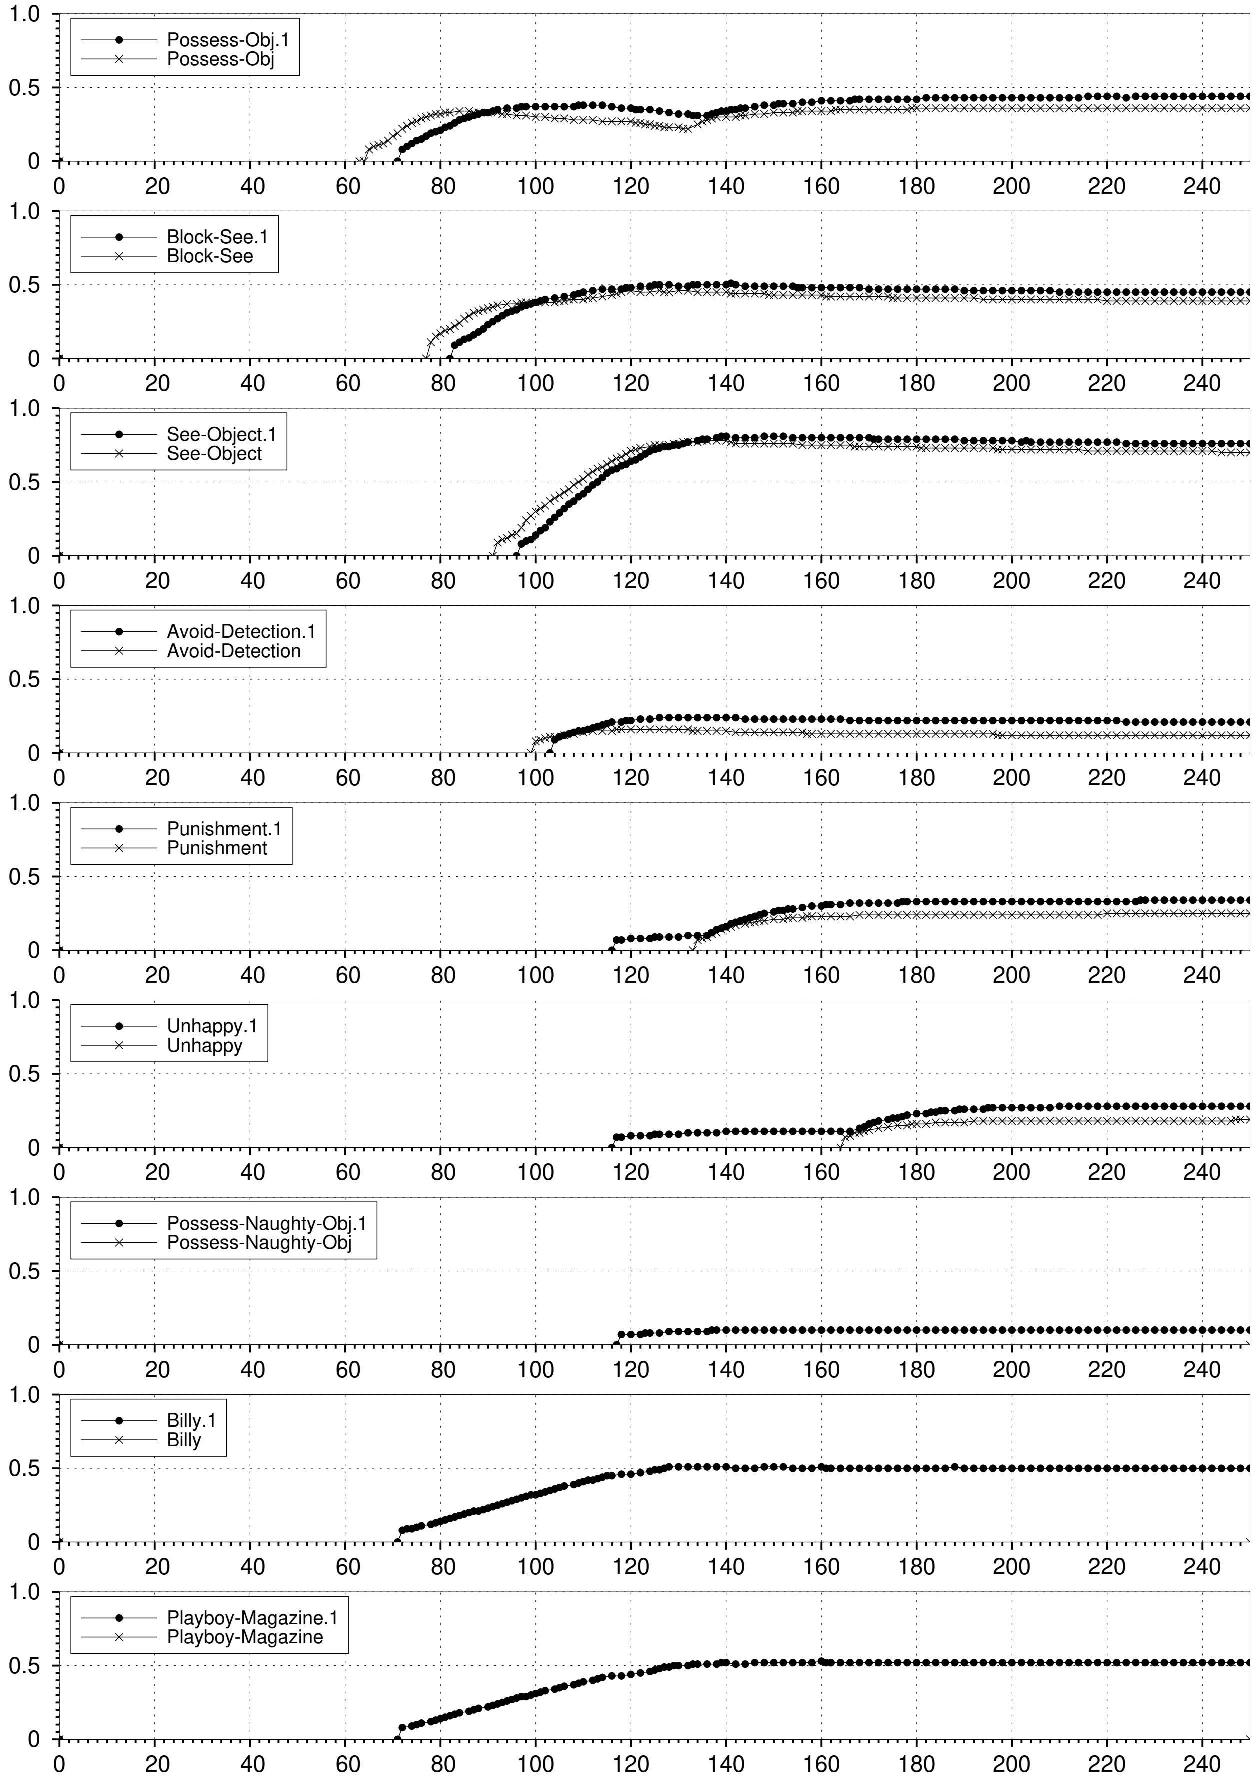

Figure 14. Evidential activation of elements of Episode.1 (**Dirty Magazine**) after presentation of **Hiding Pot**.

episodes, such as Billy.1 and Playboy-Magazine.1 in **Dirty Magazine**, do tend to become active because they play a part in so many of its roles.

5. Experiments and Discussion

REMIND has been implemented and tested in the DESCARTES connectionist simulator (Lange, Hodges, Fuenmayor, & Belyaev, 1989). The knowledge base used for the examples in this paper currently includes 206 distinct conceptual frames (e.g. Inside-Of, Avoid-Detection, Cooking-Pot) and 333 inference rules (e.g. R1 - R4). It has understood and retrieved the examples presented here (including **Hiding Pot**, **Dinner Party**, and the episodes of Table 3) and a number of other episodes of similar lengths and complexity.

In this section we discuss three simulations that illustrate (1) the importance of inferences and disambiguation on retrieval in REMIND, (2) the strong influence of superficial feature similarities on retrieval, and (3) the affect of episodic recall on the understanding process. We then compare REMIND to the ARCS and MAC/FAC models of general reminding. Finally, we discuss several directions we are exploring to extend the model.

5.1 Importance of Inferences

The retrieval of **Dirty Magazine** when **Hiding Pot** was presented as a cue to REMIND illustrates the importance of inferencing and disambiguation in the retrieval process. **Dirty Magazine** was retrieved over the more superficially similar **Flower Planting** and **Car Wash** episodes because **Dirty Magazine** and **Hiding Pot** were more thematically similar, since both involved somebody hiding something to avoid punishment. Without being able to infer these similarities, the model would not have been able to retrieve the most analogous episode.

Another example of how text comprehension affects REMIND's retrieval process is shown when two superficially-similar cues that have entirely different interpretations are presented to the network. On the surface, the cue *John put the pot inside the dishwasher because company was coming* (**Dinner Party**) is nearly the same as **Hiding Pot.** The only difference is that *company* is coming rather than *the police*. Analogical retrieval models such as ARCS (Thagard et al., 1990) and MAC/FAC (Gentner & Forbus, 1991) that do not have inferencing mechanisms would therefore have to predict that the same episode would be retrieved for both cues from amongst those shown in Table 3. This would occur for two reasons. First, the only surface feature difference between the two cues does not make a difference in overall surface similarity between the cues and any of the episodes from Table 3 — *company* and *police* are equally similar to all objects in these episodes. Second, the isomorphic structure (structural consistency) of the two cues is the same.

In contrast to ARCS and MAC/FAC, REMIND retrieves different episodes for **Dinner Party** and **Hiding Pot.** This is because REMIND interprets the two cues very differently. Whereas in **Hiding Pot** it appears John is trying to hide marijuana from the police, in **Dinner Party** it appears that he is trying to clean a cooking pot in preparation for company coming over. As such, **Dinner Party** is not likely to cause a reminding of **Dirty Magazine** and its hiding event. It seems more likely that **Dinner Party** would cause reminding of another episode involving somebody cleaning something in preparation for entertaining, such as the **Car Wash** episode (*Fred put his car in the car wash before his date with Wilma).*

This is exactly what happens in REMIND when episodic memory contains **Dirty Magazine**, **Car Wash**, and the other six episodes of Table 3. Figure 15 shows the evidential activations of the meanings of *pot* and some of the competing refinements of Inside-Of after REMIND is given input for **Dinner Party**. As when **Hiding Pot** is presented, Planting-Pot initially becomes highly activated because the input for *John put the pot inside the dishwasher* (**P1**) contains the word *pot*. After about cycle 60, REMIND infers Inside-Of-Dishwasher and Inside-Of-Opaque. Cooking-Pot's activation therefore starts to climb because of its unmatched activation from Inside-Of-Dishwasher and the other $Dishwasher-Cleaning frames. Input for *company was coming* is presented to REMIND at cycles 51-61. Spreading signature and evidential activation makes candidate inferences for *company was coming* starting from Transfer-Self, through Proximity-Of and then the Dinner-Party frames (representing somebody making dinner for somebody else). Activation eventually reaches the Clean-Dishes frame (a *precondition* for Dinner-Party). The cleaning frames were already active after the initial inferences from **P1**, so activation continues to spread and provide added support for the Inside-Of-Dishwasher refinement of Inside-Of at about cycle 130. Inside-Of-Opaque also gets added evidence at about cycle 140 from inferences through the Block-See frames from *company*

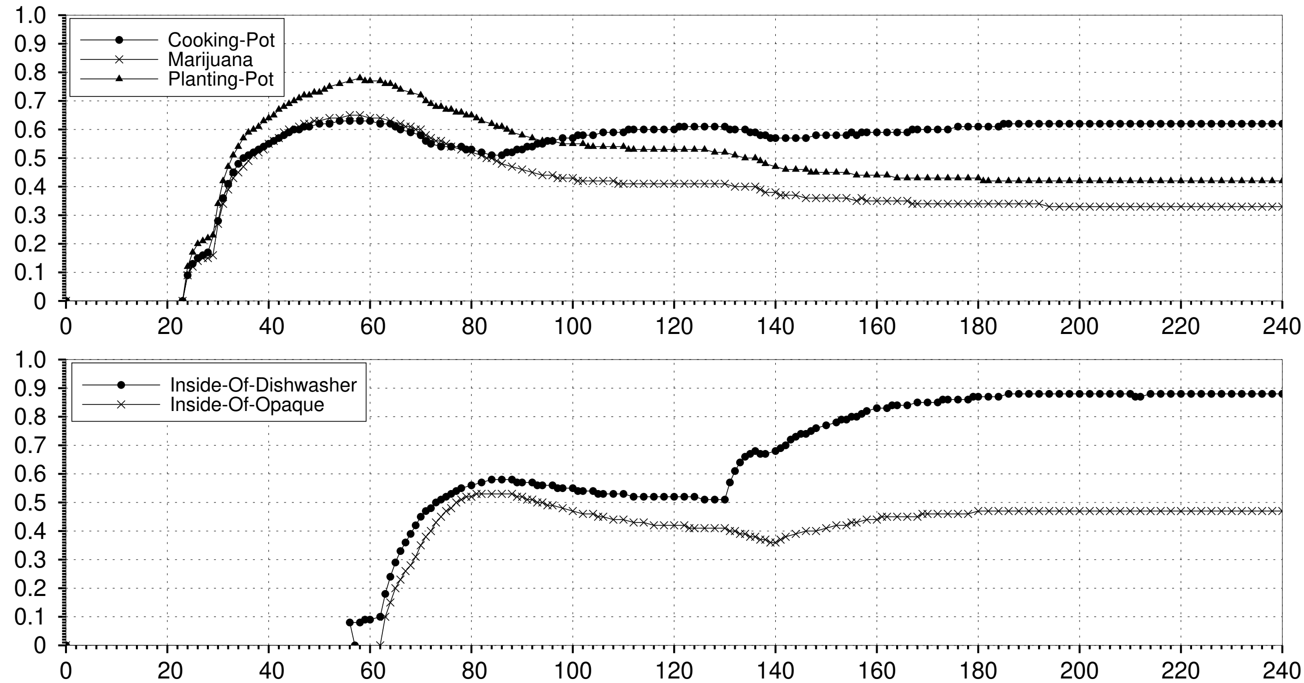

Figure 15. Evidential activations for meanings of *pot* and of competing refinements of Inside-Of after presentation of input for *John put the pot inside the dishwasher* (P1) at cycles 1 through 31 and *company was coming* (P2) at cycles 51 through 61 (**Dinner Party**)

was coming (since John could still have been hiding the pot). However, Inside-Of-Opaque does not receive enough activation to compete with Inside-Of-Dishwasher. Similarly, Marijuana never seriously competes with Cooking-Pot, because the Police-Capture and Police-See-Illegal frames that gave it unique activation in **Hiding Pot** are not inferred here (since Company does not match their selectional restrictions). The network's final interpretation of **Dinner Party** is therefore that John was trying to clean a Cooking-Pot to prepare for a Dinner-Party he was giving for Company.

Figure 16 shows the activation of the eight episodes and of two of **Car Wash**'s elements as **Dinner Party** is processed. As when **Hiding Pot** is presented, Episode.6 (Mike put the flower in the pot, and then watered it) is the first episode to become activated because of its surface similarities with *John put the pot inside the dishwasher*. Episode.2 (**Car Wash)** gets activated after about cycle 45 because of the shared inferences of Transfer-Inside and Inside-Of, and continues to gain more activation when Clean is inferred at around cycle 100. The activation of Episode.1 (**Dirty Magazine**) also starts climbing at around cycle 100 because the network is also activating the hiding and seeing structures as possible inferences for **Dinner Party**. At about cycle 120, the network infers that an Entertainment is being planned (the Dinner-Party). This activates Entertainment.2 and provides more evidence for **Car Wash**. At about the same time, the Cleaning frames (including Inside-Of-Dishwasher) get reinforcement from the *company was coming* inferences (cycles 100-140). This boosts the activations of many of the remaining elements of **Car Wash**, such as Clean.2. Accordingly, Episode.2's activation begins to dominate over Episode.1's (**Dirty Magazine**), which gradually loses support starting about cycle 150. Thus, REMIND retrieves the **Car Wash** episode when **Dinner Party** is presented as a cue. This demonstrates that even changing a single word in the cue (from *police* to *company)* can completely change the inferences and interpretations REMIND makes, and, consequently, the episode it retrieves.

5.2 Superficial Similarities

As mentioned earlier, human reminding seems to be affected strongly by superficial feature overlap between cues and memory episodes. Even when episodes exist in memory that are highly analogous to a cue (and hence useful for problem solving), people may instead get reminded of non-analogous episodes simply because they are more superficially similar. For example, stories about sharks eating people, such as **Killer Shark**, may remind people of other stories about sharks or other man-eaters devouring people, rather than thematically similar stories that do not

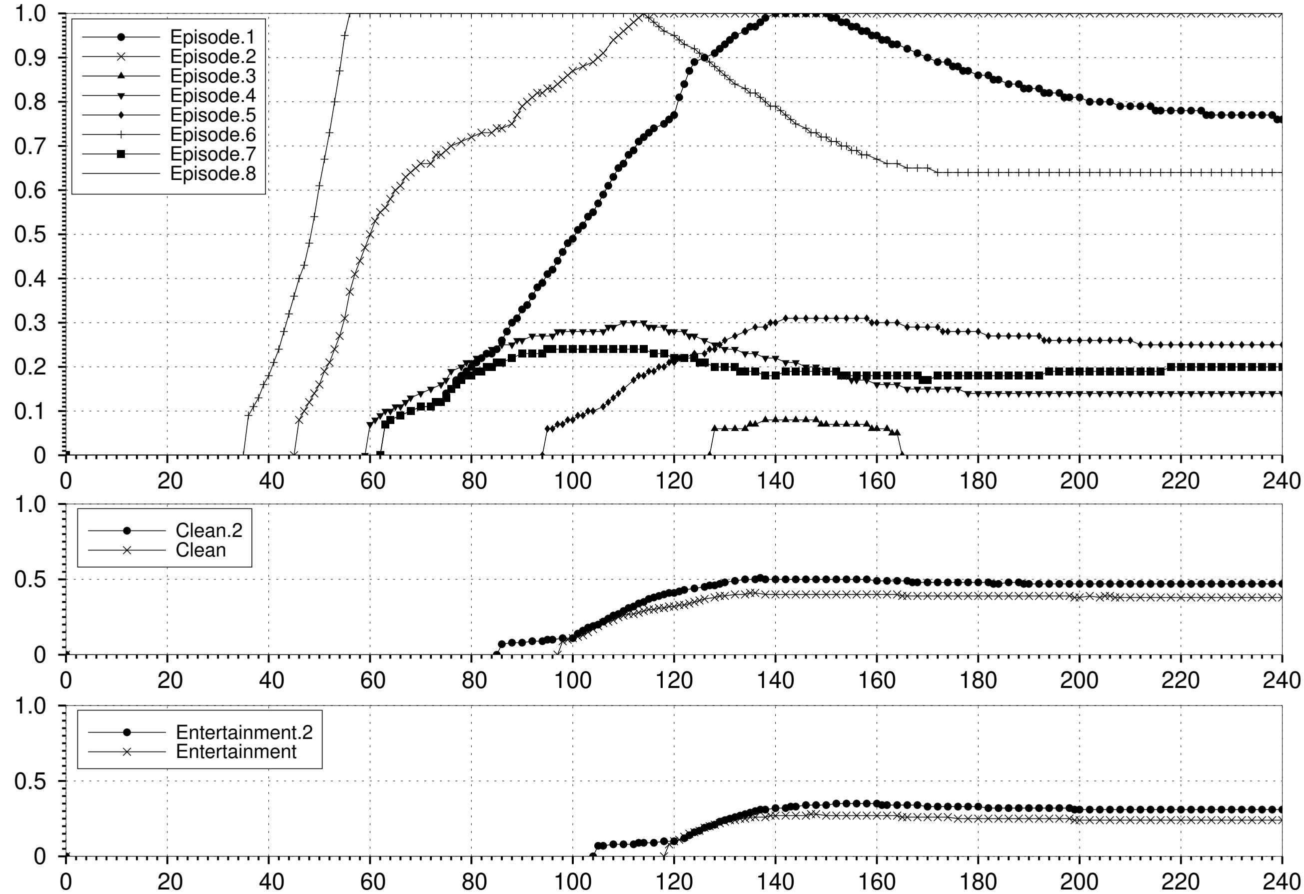

Figure 16. Activations of eight episodes of Table 3 and two of the elements of winning Episode.2 (**Car Wash**) after presentation of **Dinner Party**.

happen to involve sharks. Superficial remindings therefore often come at the expense of perhaps more valuable cross-contextual remindings.

This characteristic of reminding can be explained quite readily if episodes are remembered by storing all aspects of their interpretation, as in REMIND. Because REMIND stores all of the knowledge structures used in building an interpretation of the episode, from surface features to abstract inferences, it predicts that episodes that share many surface features with a cue are indeed likely candidates for reminding. For example, consider the sentence:

Cheech put the grass inside the bong because Chong was coming. (**Cheech and Chong**)

Cheech and Chong is an example of a superficially similar episode that can prevent retrieval of an analogous episode. Although **Cheech and Chong** is not analogous to **Hiding Pot**, the two episodes share a number of surface features. Both involve marijuana, marijuana being put inside of something, and somebody coming near. The plans and goals in the two episodes are completely different, however. In **Cheech and Chong**, the most probable interpretation is that Cheech was readying the marijuana to be smoked with his friend Chong. In **Hiding Pot**, of course, the network's interpretation was that John was hiding the marijuana so he would not be punished. So if **Cheech and Chong** is understood and stored in memory along with the eight episodes of Table 3, then **Dirty Magazine** is still the most *analogous* story in memory to **Hiding Pot**. **Cheech and Chong**, however, shares far more *total* features with **Hiding Pot**'s interpretation, and is therefore more likely to be retrieved by the model when **Hiding Pot** is presented as a cue.

To test this, input for **Cheech and Chong** was presented to the network to be understood and then remembered. The network disambiguated *grass* to Marijuana (instead of Lawn-Grass), and inferred an interpretation that Cheech put the marijuana inside a Marijuana-Bong to Light it for the Pot-Party that Chong was transferring himself to. **Cheech and Chong**'s interpretation was then added (as Episode.9) to the eight episodes already stored in the network.

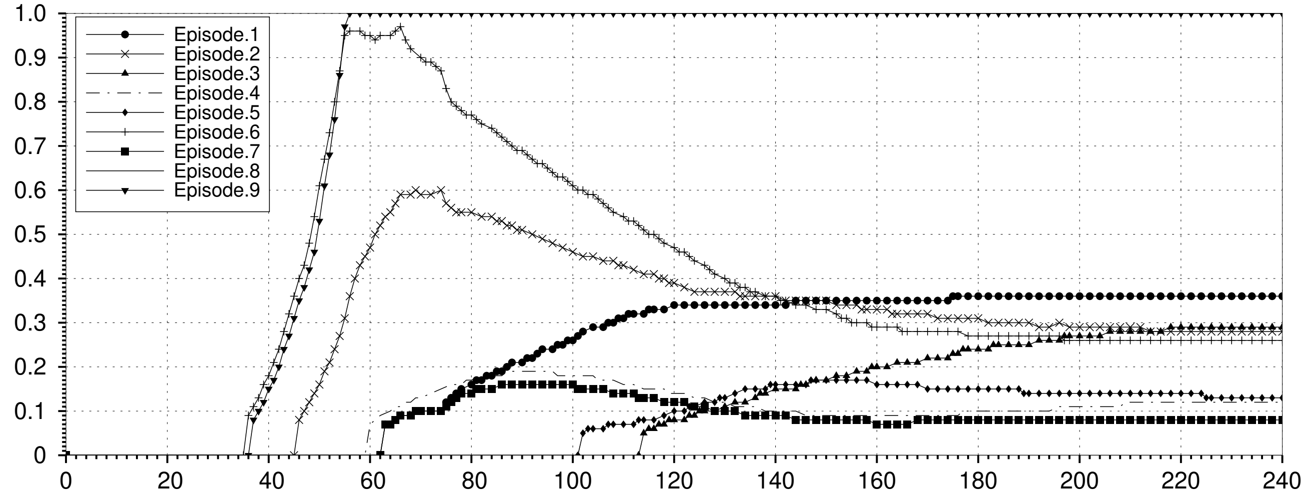

Figure 17. Activations of eight episodes of Table 3 and Episode.9 (**Cheech and Chong**) after presentation of **Hiding Pot.**

Figure 17 shows the evidential activations as **Hiding Pot** is being understood in the network containing **Cheech and Chong** and the eight other episodes of Table 3. Both **Cheech and Chong** (Episode.9) and **Flower Planting** (Episode.6) quickly become activated because of their surface similarities with the undisambiguated *John put the pot inside the dishwasher*. However, Episode.9's activation starts to dominate and Episode.6's starts to fall after *the police were coming* is presented to the network. This occurs because *the police were coming* adds more superficially similar activation to Episode.9's *Chong was coming*. Episode.6's activation drops further when Planting-Pot's activation falls as the network disambiguates *pot* to Cooking-Pot (the cleaning inferences) or Marijuana (the hiding and police capture inferences). Episode.9's activation, on the other hand, remains at its maximum (1.0) until the network settles. **Cheech and Chong** is therefore the episode retrieved for **Hiding Pot**. The other eight episodes end up with relatively little activation, since **Cheech and Chong** is so (superficially) similar to **Hiding Pot** relative to them. As might be expected, **Dirty Magazine** is the most highly activated of the remaining eight episodes, since it was the episode most similar to **Hiding Pot** before **Cheech and Chong** was remembered.

5.3 Effect of Reminding on Interpretation

REMIND only models how cues are understood and how episodes are consequently retrieved from memory. Unlike case-based reasoning models, it does not model how the information in those episodes can be used for analogical transfer or applied for problem-solving. One of the reasons this would be difficult for REMIND is that the *analogical inferences* that are made by case-based reasoning models are essentially equivalent to applying *novel* rules — e.g. applying a newly-mapped rule from a previous case to the bindings of the new case. Structured connectionist networks such as REMIND cannot currently represent such completely novel rules, because rules connecting concepts are hard-coded with units and links that cannot themselves be dynamically recruited. The possible extensions to signatures' representational and inferencing capabilities discussed in Lange (1992) might solve some of these problems. Eskridge (in press) and Barnden and Srinivas (in press) show that *hybrid* connectionist models that use complex symbolic abilities can perform analogical and case-based transfer from the cases they retrieve.

Though it does not reason from the episodes it recalls, REMIND's integration of the reminding and understanding processes shows that memory retrieval can have pragmatically interesting and useful effects on the understanding process. Episodes that become active during the understanding process feed activation back into the inferencing part of the network. This, in effect, can prime and bias the interpretation REMIND settles on for a given input. For example, consider the following example:

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

There are two possible reasons for the movie star starting to cry after the astronomer proposed to her: either the proposal made her extremely happy, or the proposal made her extremely sad. Perhaps the most likely reason for her crying was that the proposal reminded her of murdered lover, therefore making her upset and sad. REMIND, however, does not have the complex knowledge about memories and how they affect people's emotions that would be

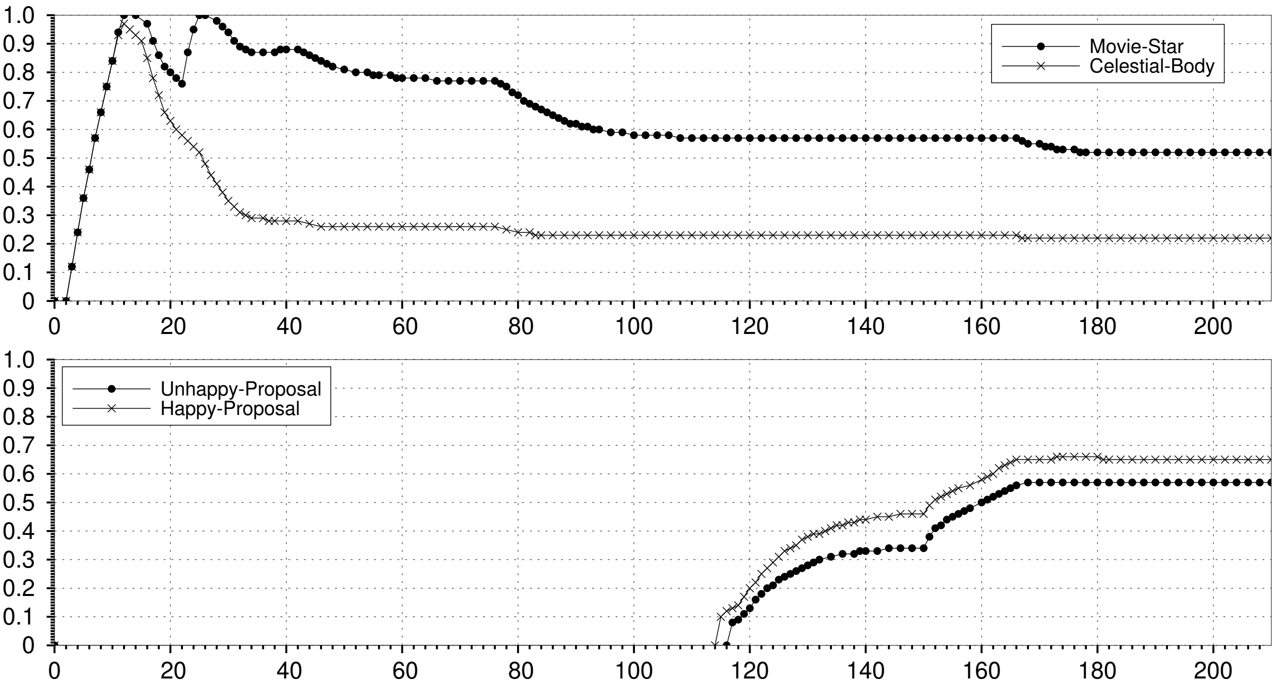

Figure 18. Activations of ambiguous meanings of word *star* and of Happy-Proposal and Unhappy-Proposal interpretations of Propose-Marriage after activation is presented for **Astronomer Proposal**.

necessary to make that interpretation. However, REMIND is given the knowledge that a person will become sad when someone they love dies (Unhappy-Dead-Friend). The network also knows that marriage proposals can be either happy events (Happy-Proposal) or sad events (Unhappy-Proposal), as in ATLAST (Eiselt, 1987).

When the network is presented with input for **Astronomer Proposal**, the word *star* is quickly disambiguated to Movie-Star because of the selectional restrictions that only Humans can be the Actors of Loves (Figure 18). REMIND then infers that the $Shooting causes the plumber to be Dead, and that the Movie-Star that Loves him will therefore be sad (Unhappy-Dead-Friend leading to Unhappy). After input for the phrase *the astronomer proposed to her* is presented, the network infers at about cycle 120 that there are two possible results from this Marriage-Proposal: that she will consider it a Happy-Proposal or an Unhappy-Proposal. These inferences then instantiate the more general Happy and Unhappy frames, respectively, both of which connect to the Movie-Star's Crying, because Crying can be the result-of states Happy or Unhappy.

As shown in Figure 18, Happy-Proposal initially becomes more highly activated than Unhappy-Proposal. This is a result of the network being biased to normally consider Marriage-Proposals to be Happy-Proposals — the weight from Marriage-Proposal to Happy-Proposal is 0.6, but the weight from Marriage-Proposal to Unhappy-Proposal is only 0.4. The gap in Happy-Proposal and Unhappy-Proposals activations begins to narrow at about cycle 150 when both begin to receive feedback from Crying (*she started to cry).* This narrowing occurs because (a) Crying has a higher weight to Unhappy than to Happy and (b) Unhappy is already highly active from Unhappy-Dead-Friend. However, REMIND's bias towards Marriage-Proposals being Happy-Proposals is too great, and Happy-Proposal finishes with more activation than Unhappy-Proposal. The final interpretation of **Astronomer Proposal** is therefore that (1) the Movie-Star was made happy by the astronomers proposal and started to cry tears of joy, and (2) she was also sad because her lover the plumber was killed (an active inference path for which REMIND could find no causal connection to the crying of the marriage proposal).

As **Astronomer Proposal** illustrates, REMIND often comes up with 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**)

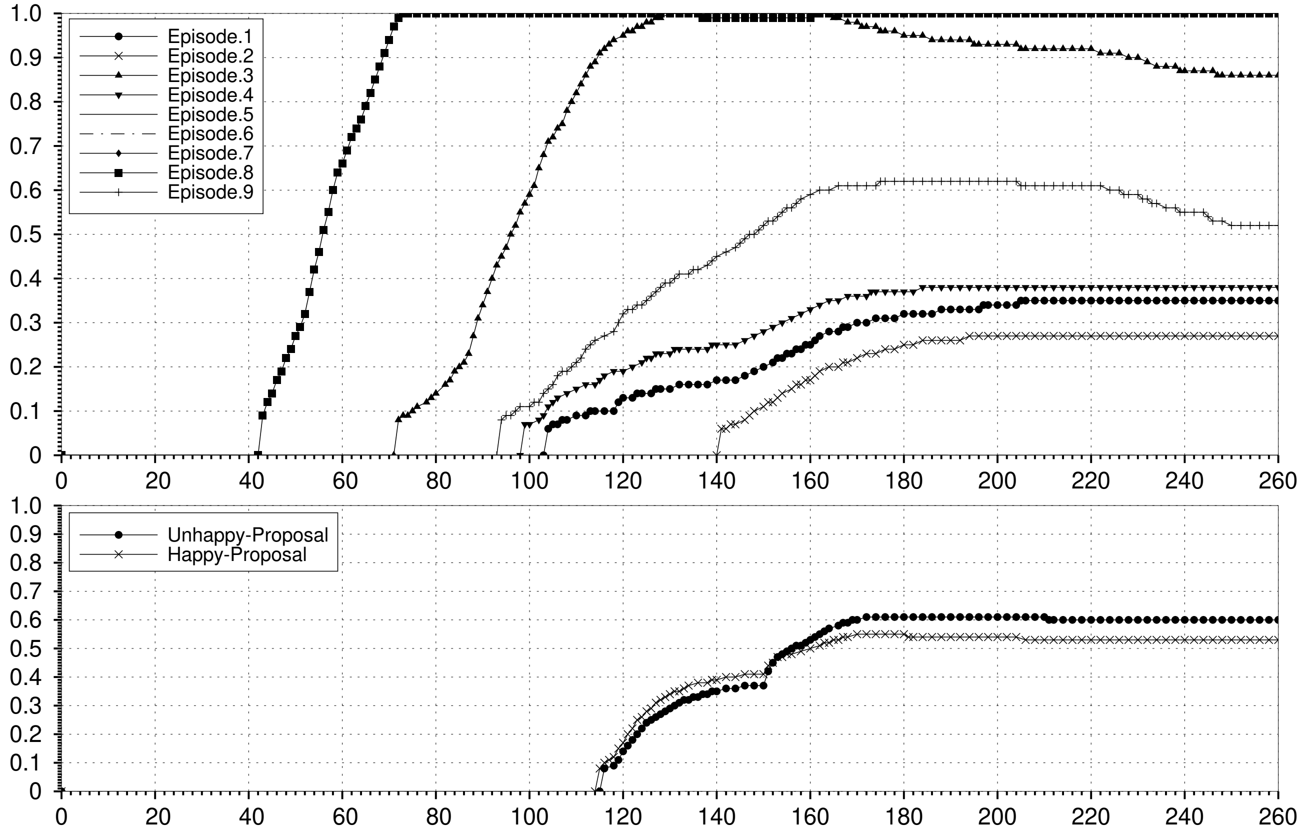

Figure 19. Activations of episodes and of Happy-Proposal and Unhappy-Proposal interpretations of Propose-Marriage after activation is presented for **Astronomer Proposal** in network containing **Sad Proposal** episode.

Sad Proposal is quite similar to **Astronomer Proposal**, except that in **Sad Proposal** the input made it explicit that Suzie became Unhappy after Bill proposed to her. This essentially forces 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 Table 3.

Figure 19 shows the activation levels of **Sad Proposal** (Episode.8) and the other episodes as **Astronomer Proposal** is being understood by REMIND. As expected, **Sad Proposal** quickly dominates most of the other episodes because it is so similar to **Astronomer Proposal**. Episode.3 becomes temporarily active because it also involves somebody shooting somebody to death. However, **Sad Proposal** eventually wins and is retrieved.

The most interesting result shown in Figure 19 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 because of its higher weight from Propose-Marriage. In this case, however, Episode.8 is highly active, and with it Unhappy.8 and Unhappy-Proposal.8. As described in section 4, these episode instances feed activation back into their concepts in the understanding network. Unhappy-Proposal therefore gets significant activation from Unhappy-Proposal.8. This added evidence allows its activation to climb over that of Happy-Proposal, which gets no added evidence from any of the episodes in memory. When the network settles, Unhappy-Proposal wins over Happy-Proposal, so REMIND's interpretation is that the astronomer's marriage 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 the same spreading-activation mechanism for both language understanding and episodic memory retrieval demonstrates one way memory retrieval can subtly affect the interpretation process. When stored episodes share conceptual similarity with a cue that REMIND is comprehending, these episodes feed evidential 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. These effects emerge entirely from the integration of language understanding and retrieval within a single spreading-activation network.

5.4 Comparison to General Models of Reminding

It is difficult to directly compare REMIND to most case-based reasoning models because they were developed with different goals in mind. As described previously, CBR models are usually models of expert reminding or models built to demonstrate how certain kinds of abstract remindings can occur. 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 allow 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.

The models that REMIND is most directly comparable to are ARCS (Thagard et al., 1990) and MAC/FAC (Gentner & Forbus, 1991), two other simulations of general, non-expert reminding. All three models were built to take into account psychological evidence showing that episodic memory retrieval is strongly influenced by surface feature similarities between a cue and the target episodes in memory, but that deeper analogical or thematic similarities sometimes play an important role. However, there are two important differences between how REMIND explains this evidence compared to ARCS and MAC/FAC.

The most obvious difference is that REMIND is an inferencing-based theory of reminding. Memory retrieval in REMIND results directly from the inferencing and disambiguation process used to understand and form elaborated interpretations of REMIND's cue. ARCS and MAC/FAC, on the other hand, are stand-alone memory retrieval models that are given complete propositional representations of their cues and memory episodes. An advantage of this approach is that it allows ARCS and MAC/FAC to deal with retrieval of much more complicated episodes than does REMIND. ARCS, for example, performed memory retrieval using complex, hand-coded predicate calculate representations of synopses of 24 Shakespearean plays. On the other hand, a major criticism of ARCS and MAC/FAC is that neither model specifies exactly how the representation of its input cues and episodes is formed. More importantly, neither model specifies what kinds of knowledge those representations should generally include. Should the cue representations include only the surface propositions directly stated in phrases of 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?

We believe that many types of memory retrieval cannot be performed without such inferences, as illustrated by some of the examples in this paper (e.g., section 4.3 and 5.1) and some of the examples of case-based reasoning systems. However, even if stand-alone retrieval models such as ARCS and MAC/FAC used fully elaborated interpretations of their cues, we believe that not modeling the *proces*s by which these interpretations are formed misses an important factor in reminding. People read texts with widely different levels of analysis, ranging from simply skimming the text to reading it carefully and thinking deeply about its ramifications. While there are circumstances under which it is reasonable to assume that subjects make relatively deep thematic inferences, it is misleading to think that this is always the case (e.g., Seifert et al., 1986). Thus, when the process by which the retrieval cue is constructed itself 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. In contrast, REMIND explicitly models the cue interpretation process, and so can potentially explain when elaborated abstract inferences are available to affect reminding.

Another major difference between REMIND and ARCS and MAC/FAC is in how they theorize that analogical similarity exerts an influence on memory retrieval. Both ARCS and MAC/FAC perform memory retrieval in two stages. In their first stage, both search memory to find the episodes in memory sharing the most surface semantic commonalties with the cue. In their second stage, they compute which of the contacted episodes best match the cue and should be retrieved. In addition to counting surface similarities, they calculate the degree of *structural isomorphism* (or *relational consistency)* as an explicit factor in their computation of which episode best matches the cue. Eskridge (in press) describes a hybrid connectionist model of analogical retrieval and transfer that acts in a similar way.

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 therefore more isomorphic, in a purely syntactic sense, to **Boy Run** than is **Rover Run**, and is therefore a better analog.

Both ARCS and MAC/FAC explicitly compute the level of syntactic isomorphism between a cue and episodes that share surface semantic overlap with the cue. The degree of cue/target isomorphic match is factored into their second stages matching process. Analogical similarity is hypothesized by these models to exert its effect on memory retrieval as a direct result of this specifically computed degree of syntactic relational consistency between cues and memory episodes. REMIND, on the other hand, never explicitly computes the degree of isomorphism between a cue and memory episodes. In REMIND, the influence of such relational consistency is entirely the result of the inferencing process. Relationally consistent targets are retrieved over relationally inconsistent targets in REMIND only 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 in the form of anger or a kick from the boy. Because of the different inferences and interpretation of the two episodes, REMIND, like ARCS and MAC/FAC, would therefore 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.

Like ARCS, our earlier model of integrated language understanding and memory retrieval, SAARCS (Lange, Melz, Wharton, and Holyoak, 1990), included relational consistency as an explicit factor in reminding. SAARCS was a hybrid connectionist model that combined ROBIN (Lange & Dyer, 1989) with aspects of ARCS. Like REMIND, SAARCS used the ROBIN portion of the network to infer and disambiguate an interpretation of a cue. Unlike REMIND, SAARCS then explicitly calculated relational consistency to build a constraint satisfaction mapping network like ARCS' to determine which episode was retrieved. Oftentimes SAARCS would not have enough knowledge to make different inferences between cues with cross-mapped bindings, such as for *The boat followed the dolphins* vs. *The dolphins followed the boat*. In those cases, SAARCS' use of ARCS' explicit influence of relational consistency between cues and targets led SAARCS to retrieve the right analog (Lange et al., 1990). However, when SAARCS had enough knowledge to build a different interpretation for cross-mapped cues (such as for *The surfer ate the shark* vs. The *shark ate the surfer)*, the explicit effect of relational consistency turned out to be unneeded. As in REMIND, the different shared inferences in SAARCS were enough to activate the analogous episode enough to win. It turned out that in every case that we built enough knowledge into the network for it to build different interpretations for disanalogous episodes, the degree of syntactic isomorphism only boosted (confirmed) the analogous episode that had already won due to those different interpretations. This was a primary motivating factor for simplifying the model and moving to REMIND, a purely inferencing-based model of episodic reminding.

We therefore believe that the noted effects of syntactic isomorphism and relational consistency on memory retrieval can be fully explained by the understanding process. Relationally consistent episodes tend to have similar inferences, interpretations, and themes, while relationally inconsistent episodes tend to have different inferences, interpretations, and themes. In REMIND, relationally-based reminding occurs as a natural side-effect of interpreting and disambiguating an input text. Relational consistency only affects reminding to the degree that it changes those inferences.

5.5 Future Work

In the future, there are three main areas that we would like to explore: (1) improved inferencing abilities, (2) the ability to determine the initial surface role-bindings with additional lexical information in the networks, and (3) automatic learning of the episodes the network has understood.

5.5.1 Improved Inferencing

REMIND's propagation of signature activations dynamically instantiates candidate inference paths in parallel, in much the same way as marker-passing systems. The use of ROBIN's signatures therefore allows REMIND to perform dynamic inferencing difficult for most connectionist models, while using its evidential activation allows it to perform disambiguation and reinterpretation difficult for most symbolic models. However, REMIND's representation and rule-firing abilities are currently limited relative to those of traditional symbolic models, limiting the length and complexity of the texts the model can understand and therefore remember.

One of the main restrictions of the model as described is that there can be only one dynamic instance of each frame at any given time, since binding units can only hold one signature activation at once. Because of this, REMIND cannot yet represent or interpret texts involving two different seeing or eating events, for instance. Another limitation is that REMIND currently only propagates signatures of *pre-existing* concepts, such as of Cooking-Pot, Marijuana, or John. REMIND does not propagate signatures of the *dynamically* instantiated frames inferred by signatures, such as the dynamic instance of Cooking-Pot or Marijuana being Inside-Of a Dishwasher in Figure 5. Not being able to propagate signatures of dynamically instantiated frames makes it impossible for REMIND to encode most rules for *general* planning knowledge or complex interactions of goals, which generally require the ability to reason over any dynamic plan or goal instance the system might have. These type of rules are needed to understand many complex texts, such as those involving abstract planning failures or themes (cf. Schank, 1982; Dyer, 1983).

These and other restrictions on the types of inferencing REMIND performs limit the complexity of the episodes REMIND can currently understand and retrieve relative to many symbolic language understanding and case-based reasoning models. We are currently exploring solutions to some of these problems to allow multiple dynamic instantiations of individual frames and to allow rules that propagate signatures of the novel concepts inferred by the network. A number of different ways to approach these problems and handle more complex text are discussed in Lange (1992). Shastri and Ajjanagadde (in press) discuss an analogous solutions for structured networks that do not perform disambiguation.

5.5.2 Lexical Information and Initial Role-Bindings

REMIND does not currently address the problem of deciding upon the original syntactic bindings, e.g. that *pot* is bound to the Object role of a phrase. Rather, REMIND's networks are given these initial bindings and use them for high-level inferencing. To handle natural language input entered as text, the network must somehow contain and use syntactic and phrasal information to create the initial role-bindings that REMIND is currently given by hand.

5.5.3 Automatic Learning of Episode Units

The representations that REMIND uses for its memory episodes are created entirely by REMIND's spreading-activation understanding process. To store those representations in the network, however, the units and connections used to encode them must be added by hand. It would be desirable to have those episode units be learned automatically by the network itself. We are currently exploring a mechanism to *recruit* units (cf. Diederich, 1990) to encode the interpretation held by signature and evidential activation in the network and therefore allow episode learning to proceed without intervention.

Another area to explore will be gradual decay of the weights between episode units and the semantic network. Episodes that are not retrieved for a long time should gradually have their weights decay so that they become more difficult to become reminded of as time goes on. This is also a potential solution to the problem of indexing too many episodes under general concepts, such as Possess-Obj and Inside-Of, since connections from them to old episodes could gradually decay away and become available for new ones.

6. Conclusions

The process of memory retrieval has generally been explored in isolation from the process of language understanding, even though both activities seem to share many of the same processes. One of the reasons for this is that building an integrated model of language understanding and retrieval requires not only solving the important aspects of each of the processes individually — a difficult enough problem separately — but also finding a parsimonious way to integrate the two processes that can explain the effect they have on each other. Because of their simple processing mechanism and demonstrated abilities to perform disambiguation and model the psychological effects of priming and context, structured spreading-activation networks seem to have a great deal of potential for building such an integrated model.

A number of difficult problems for spreading-activation networks have to be solved to use them for language understanding and memory retrieval. First, language understanding requires the ability to represent variable bindings within the network and to perform dynamic inferencing to explain the plans and goals of actors. Second, problems of disambiguation and frame selection require the ability to combine evidential domain knowledge with the contextual evidence of the network's inferences to select the best interpretation of a text from multiple possible interpretations. In addition, ambiguous texts requiring many inferences make the problems of controlling crosstalk between unrelated concepts and inferences especially important to solve. And finally, a method for representing episodes in the network's long-term memory that allows them to be influenced and retrieved by the spreading-activation understanding process must be developed.

This paper has described REMIND, a structured spreading-activation model that solves many of these problems in an integrated model of language understanding and episodic reminding. In REMIND, activation is spread through a semantic network that performs dynamic inferencing and disambiguation to infer a conceptual representation of an input cue. Because stored episodes are associated with concepts used to understand them, the spreading-activation process also activates any memory episodes in the network that share features or knowledge structures with the cue. After the cue's conceptual representation is formed, the network recalls the memory episode having the highest activation. Since the inferences made from a cue often include actors plans and goals only implied in a cue's text, REMIND is able to get abstract, analogical remindings that would not be possible without an integrated understanding and retrieval model.

Theoretically, REMIND lies somewhere between case-based reasoning models and general analogical retrieval models such as ARCS and MAC/FAC. Like ARCS and MAC/FAC, REMIND is meant to be a psychologically-plausible model of general human reminding, and therefore takes into account the prevalence of superficial feature similarities in remindings. However, we believe that many of the types of high-level planning and thematic knowledge structures used as indices in case-based reasoning systems also have an important effect on reminding. REMIND is thus partially an attempt to bridge the gap between case-based and analogical retrieval models. As it turns out, this gap is naturally bridged when the same spreading-activation mechanism is used to both understand cues and to retrieve episodes from memory. Using the same mechanism for both processes causes retrieval to be affected by all levels that a text was understood with. This is the case in REMIND, in which the understanding mechanism is given the superficial features and actions of a text and attempts to explain them by inferring the plans and goals being used causing long-term memory episodes to be activated by both.

Although being an integrated comprehension and retrieval model makes REMIND more complex than the retrieval mechanisms of case-based reasoning models and ARCS and MAC/FAC, it also allows REMIND to be simpler than them in a couple of significant respects. A large amount of the research in case-based reasoning models is devoted towards discovering the best indices to store cases under (the *indexing problem)*. In REMIND, however, episodes are simply stored (indexed) under all of the concepts that played a part in understanding them. When combined with the comprehension part of the model, this simple connectionist approach to indexing avoids the indexing problem altogether, while still giving the effect of having chosen the proper indices since the most salient and unique features of an episode in a given context naturally become highly activated as part of the understanding process.

ARCS and MAC/FAC also avoid the indexing problem of CBR models because they both make contact with episodes that share any feature similarities with the cue, in effect using all features as indices. Where REMIND differs from ARCS and MAC/FAC is that both ARCS and MAC/FAC use separate mechanisms to explicitly factor the degree of syntactic relational consistency (analogical similarity) into retrieval. We believe that psychological effects of analogical similarity on memory retrieval that their separate syntactic mechanisms are meant to model can be fully explained by the understanding process. Relationally consistent episodes tend to have similar inferences, interpretations, and themes, while relationally inconsistent episodes do not. In REMIND, analogically-based reminding therefore occurs as a natural side-effect of understanding an input text, rather than as a result of a separate process that explicitly computes it, as in ARCS and MAC/FAC.

A final aspect to note about REMIND concerns how language understanding and retrieval processes come full circle. The episode retrieved depends crucially on the interpretation of the cue from the spreading-activation network's inferences. Once an episode is retrieved, it in turn primes the activation of the evidential spreading-activation network, perhaps leading to a different disambiguation and therefore interpretation of the next cue. Thus, we believe that REMIND is able to uniquely provide insights which other current reminding models are not able to show. As such, REMIND represents an entirely new class of reminding models.

6.1 Acknowledgments

This research was supported by NSF grant DIR-9024251, Army Research Institute contract MDA 903-89-K-0179, and by a grant from the Keck Foundation. We would like to thank John Barnden, Michael Dyer, Keith Holyoak, Eric Melz, and John Reeves for their helpful comments on the model and on earlier drafts of this paper.

7. References

Barnden, J. (1990). The power of some unusual connectionist data-structuring techniques. In J. A. Barnden and J. B. Pollack (Eds.), *Advances in connectionist and neural computation theory*, Norwood, NJ: Ablex.

Barnden, J., & Srinivas, K. (in press). Overcoming rule-based rigidity and connectionist limitations through massively-parallel case-based reasoning. *International Journal of Man-Machine Studies*.

Charniak, E. (1986). A neat theory of marker passing. *Proceedings of the Fifth National Conference on Artificial Intelligence*, Philadelphia, PA.

Cottrell, G., & Small, S. (1982). A connectionist scheme for modeling word-sense disambiguation. *Cognition and Brain Theory, 6*, 89-120.

Dolan, C. P., & Smolensky, P. (1989). Tensor product production system: A modular architecture and representation. *Connection Science, 1 (1)*, 53-68.

Dyer, M. G. (1983). *In-depth understanding: A computer model of integrated processing for narrative comprehension*. Cambridge, MA: MIT Press.

Eskridge, T. C. (in press). A hybrid model of continuous analogical reasoning. To appear in J. Barnden & K. Holyoak (Eds.), *Advances in connectionist and neural computation theory, volume II: Analogical connections*. Norwood, NJ: Ablex.

Falkenhainer, B., Forbus, K. D., & Gentner, D. (1989). The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, *41*, 1-63.

Gentner, D. (1989). The mechanisms of analogical learning. In S. Vosniadou and A. Ortony (Eds.), *Similarity and analogical reasoning,* (pp. 199-241). New York, NY: Cambridge University Press.

Gentner, D., & Forbus, K. D. (1991). MAC/FAC: A model of similarity-based retrieval. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 504-509). Hillsdale, NJ: Lawrence Erlbaum.

Gentner, D., & Landers, R. (1985). Analogical reminding: A good match is hard to find. *Proceedings of the International Conference on Systems, Man and Cybernetics* (pp. 607-613). Tucson, AZ.

Gick, M., & Holyoak, K. J. (1980). Analogical problem solving. *Cognitive Psychology, 12*, 306-355.

Gick, M. L., & McGarry, S. J. (in press). Learning from mistakes: Inducing analogous failures to a source problem produces later successes in analogical transfer. *Journal of Experimental Psychology: Learning, Memory, and Cognition.*

Glucksberg, S., Kreuz, R. J., & Rho, S. H. (1986). Context can constrain lexical access: Implications for models of language comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 12,* 323-335.

Granger, R. H., Eiselt, K. P., & Holbrook, J. K. (1986). Parsing with parallelism: A spreading activation model of inference processing during text understanding. In J. Kolodner and C. Riesbeck (Eds.), *Experience, memory, and reasoning* (pp. 227-246). Hillsdale, NJ: Lawrence Erlbaum.

Hammond, K. (1989) *Case-based planning.* Boston: Academic Press.

Hendler, J. (1989). Marker-passing over microfeatures: towards a hybrid symbolic/connectionist model. *Cognitive Science, 13,* 79-106.

Hofstadter, D., & Mitchell, M. (in press). The copycat project: A model of mental fluidity and analogy-making. To appear in J. Barnden & K. Holyoak (Eds.), *Advances in connectionist and neural computation theory, volume II: Analogical connections*. Norwood, NJ: Ablex.

Holldobler, S. (1990). A structured connectionist unification algorithm. *Proceedings of the Ninth National Conference on Artificial Intelligence*, Boston, MA.

Holyoak, K. J., & Koh, K. (1987). Surface and structural similarity in analogical transfer. *Memory & Cognition*, *15*, 332-340.

Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review, 95*, 163-182.

Kitano, H., Tomabechi, H., & Levin, L. (1989). Ambiguity resolution in DMTrans Plus. *Proceedings of the Fourth Conference of the European Chapter of the Association of Computational Linguistics*. Manchester University Press.

Kohonen, T. (1984). *Self-organization and associative memory*. New York: Springer Verlag.

Kolodner, J. (1984). *Retrieval and organizational strategies in conceptual memory: A Computer Model*. Hillsdale, NJ: Lawrence Erlbaum.

Kolodner, J., Simpson, R., & Sycara, K. (1985). A process model of case-based reasoning in problem solving. *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, Los Angeles, CA.

Lange, T. (1992). Lexical and pragmatic disambiguation and reinterpretation in connectionist networks. *International Journal of Man-Machine Studies, 36,* 191-220.

Lange, T., & Dyer, M. G. (1989). High-level inferencing in a connectionist network. *Connection Science, 1 (2)*, 181-217.

Lange, T., Hodges, J., Fuenmayor, M., & Belyaev, L. (1989): DESCARTES: Development Environment For Simulating Hybrid Connectionist Architectures. In *Proceedings of the Eleventh Annual Meeting of the Cognitive Science Society* (pp. 698-705). Hillsdale, NJ: Lawrence Erlbaum.

Lange, T., Melz, E., Wharton, C., & Holyoak, K. (1990). Analogical retrieval within a hybrid spreading-activation network. In D. S. Touretzky, J. L. Elman, T. J. Sejnowski, and G. E. Hinton (eds.), *Proceedings of the 1990 Connectionist Models Summer School*. San Mateo, CA: Morgan Kaufmann.

Lebowitz (1980). *Generalization and memory in an integrated understanding system.* Ph.D. thesis, Research Report #186, Yale University, Department of Computer Science, New Haven, CT.

Lytinen, S. (1984). *The organization of knowledge in a multi-lingual integrated parser.* Ph.D. thesis, Research Report #340, Yale University, Department of Computer Science, New Haven, CT.

McClelland, J. L. & Kawamoto, A. H. (1986): Mechanisms of Sentence Processing: Assigning Roles to Constituents of Sentences. In McClelland & Rumelhart (eds.), *Parallel Distributed Processing: Vol. 2*, p. 272-325. Cambridge, MA: The MIT Press.

Miikkulainen, R. (in press). *Towards subsymbolic natural language processing: An integrated model of scripts, lexicon, and memory.* Cambridge, MA: The MIT Press.

Miikkulainen, R. & Dyer, M. G. (1991). Natural language processing with modular PDP networks and distributed lexicon. *Cognitive Science, 15*, 343-399.

Norvig, P. (1989). Marker passing as a weak method for text inferencing. *Cognitive Science, 13,* 569-620.

Owens, C. (1989). Integrating feature extraction and memory search. In *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society,* Ann Arbor, MI.

Ratterman, M. J., & Gentner, D. (1987). Analogy and similarity: Determinants of accessibility and inferential soundness. *Proceedings of the Ninth Annual Meeting of the Cognitive Science Society* (pp. 22-34). Hillsdale, NJ: Lawrence Erlbaum.

Reeves, J. F. (1991). *Computational morality: A process model of belief conflict and resolution for story understanding*. Ph.D. thesis, Computer Science Department, University of California, Los Angeles.

Riesbeck, C. K., & Schank, R. C. (1989). *Inside case-based reasoning.* Hillsdale, NJ: Lawrence Erlbaum.

Riesbeck, C. K., & Martin, C. E. (1986). Direct memory access parsing. In J. Kolodner & C. Riesbeck (Eds.), *Experience, Memory, and Reasoning*, pp. 209-226. Hillsdale, NJ: Lawrence Erlbaum.

Ross, B. H. (1989). Distinguishing types of superficial similarities: Different effects on the access and use of earlier problems. *Journal of Experimental Psychology: Learning, Memory, and Cognition,* 15, 456-468.

Rumelhart, D. E., Hinton, G. E., & McClelland, J. L. (1986): A General Framework for Parallel Distributed Processing. In Rumelhart & McClelland (eds.), *Parallel Distributed Processing: Vol. 1*, p. 45-76. Cambridge, MA: The MIT Press.

Schank, R. C. (1982). *Dynamic memory.* New York: Cambridge University Press.

Schank, R. C., & Abelson, R. (1977). *Scripts, plans, goals and understanding*. Hillsdale, NJ: Lawrence Erlbaum.

Schank, R. C., & Leake, D. B. (1989). Creativity and learning in a case-based explainer. *Artificial Intelligence, 40,* 353-385.

Seifert, C. M. (in press). The role of goals in retrieving cases. To appear in J. Barnden & K. Holyoak (Eds.), *Advances in connectionist and neural computation theory, volume II: Analogical connections*. Norwood, NJ: Ablex.

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

Shastri, L., & Ajjanagadde, V. (in press). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. To appear in *Behavioral and Brain Sciences*.

St. John, M. (in press). The story gestalt: A model of knowledge intensive processes in text comprehension. *Cognitive Science*.

Sun, R. (in press). Beyond associative memories: Logics and variables in connectionist models. To appear in *Information Sciences*.

Swinney, D. A. (1979). Lexical access during sentence comprehension: (Re)consideration of context effects. *Journal of Verbal Learning and Verbal Behavior, 15,* 681-689.

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

Till, R. E., Mross, E. F., & Kintsch, W. (1988). Time course of priming for associate and inference words in a discourse context. *Memory & Cognition, 16,* 283-298.

Touretzky, D. (1990). Connectionism and compositional semantics. In J. A. Barnden & J. B. Pollack (Eds.), *Advances in connectionist and neural computation theory*, Norwood, NJ: Ablex.

Touretzky, D., & Hinton, G. (1988). A distributed connectionist production system. *Cognitive Science*, *12*, 423-466.

Waltz, D., & Pollack, J. (1985). Massively parallel parsing: A strongly interactive model of natural language interpretation. *Cognitive Science, 9*, 51-74.

Wilensky (1983). *Planning and Understanding*. Reading, MA: Addison-Wesley.

Wharton, C. M, Holyoak, K. J., Downing, P. E., Lange, T. E., & Wickens, T. D. (1991). Retrieval competition in memory for analogies. *Proceedings of the Thirteenth Annual Conference of the Cognitive Science Society* (pp. 528-533). Hillsdale, NJ: Lawrence Erlbaum.

Wharton, C. M, Holyoak, K. J., Downing, P. E., Lange, T. E., & Wickens, T. D. (1992). The story with reminding: Memory retrieval is influenced by analogical similarity. *Proceedings of the Fourteenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.