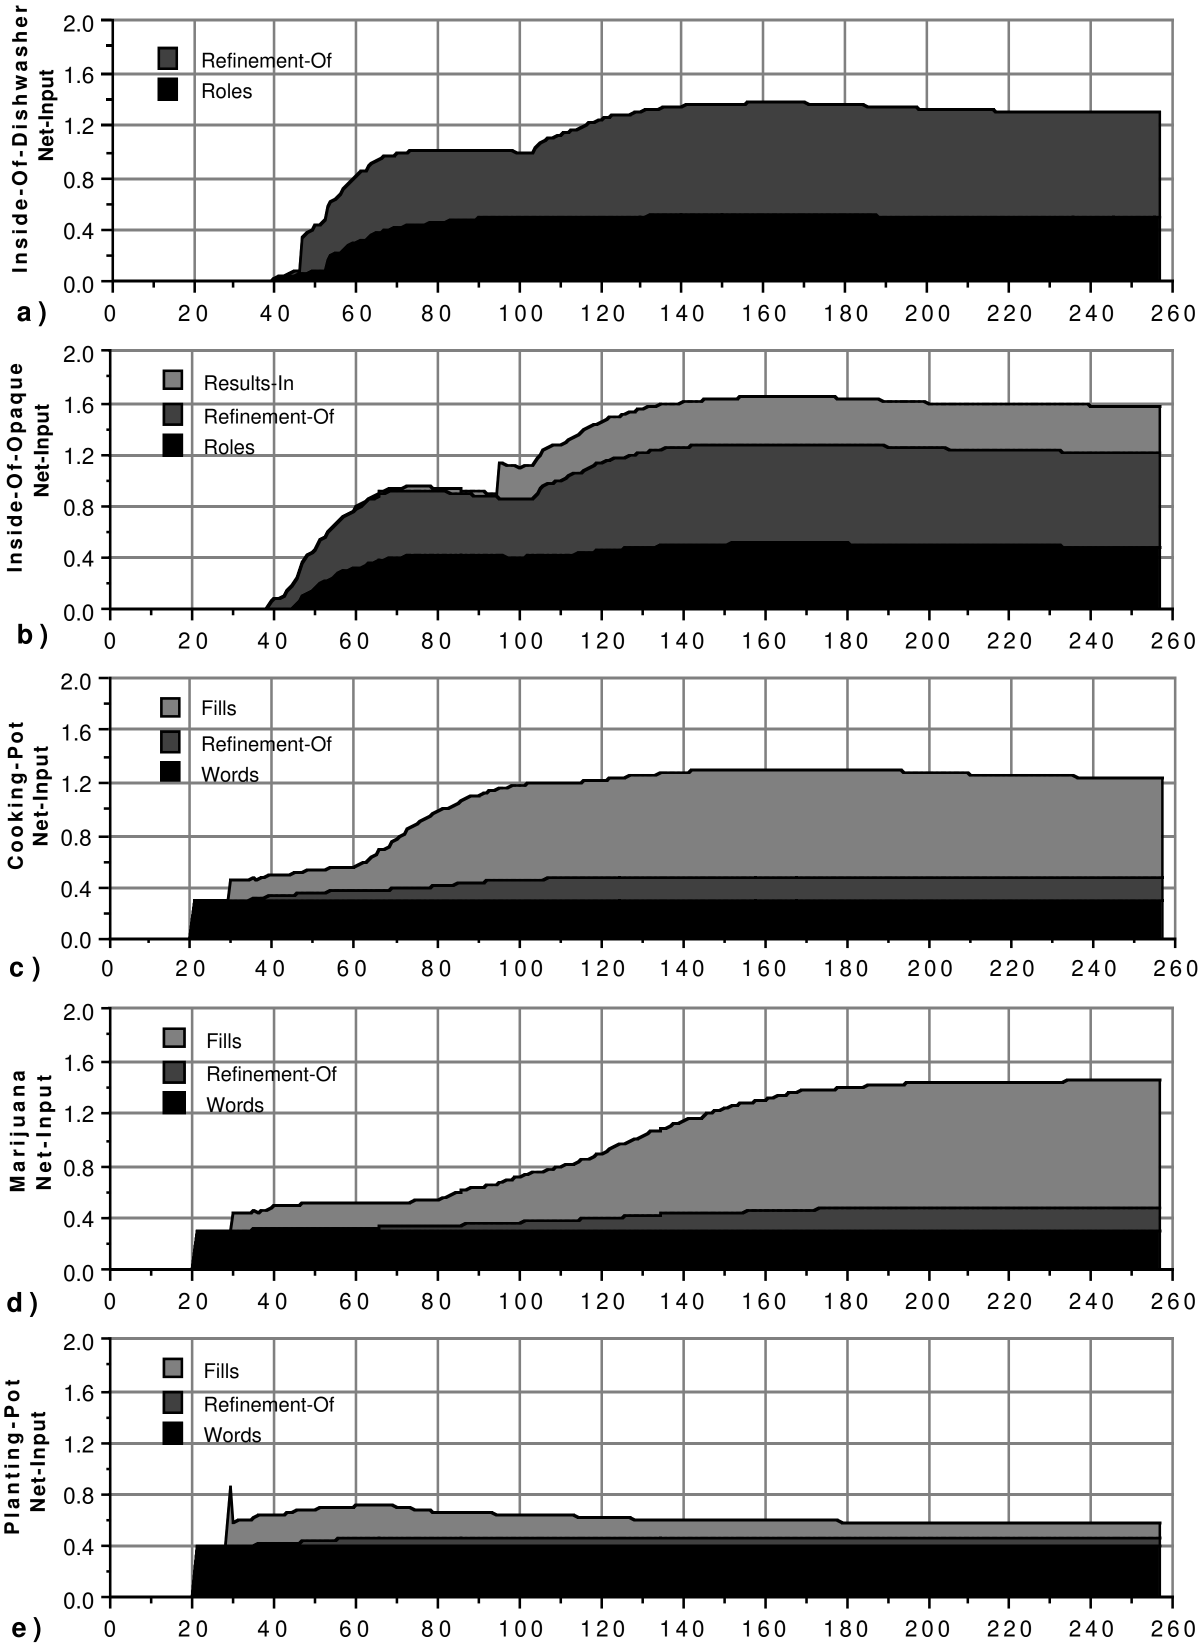

Figure 10. Activation of input branches for the evidential nodes of ambiguous frames in **Hiding Pot.**

activation from its other role fillers. The advantage of Inside-Of-Dishwasher's Role input branch over Inside-Of-Opaque's between cycles 65 and 100 is thus relatively small (Figures 10a & b), since Inside-Of-Dishwasher and Inside-Of-Opaque share most of the same role-bindings (both were planned by John, have the Location of a Dishwasher, and have Cooking-Pot as a potential Object).

While this is occurring, input for **Hiding Pot** continues to be presented to the network. At cycle 50, the lexical node *police* is clamped, as is its signature binding and the evidential node for phrase <S "were coming"> at cycle 60. Inferences are made by propagation of signature and evidential activation through Transfer-Self, Proximity-Of, and See-Object, until two things happen. First, Police-See-Illegal gets instantiated from See-Object (which has itself received evidential activation through inferences from both phrases of the sentence). Police-See-Illegal accepts only objects that are Illegal-Objects, so the network's selectional restrictions allow only Marijuana through, giving it an opportunity to receive unshared activation from Police-See-Illegal and eventually Police-Capture. Thus the activation of Marijuana's Fills input branch increases at about cycle 90 (Figure 10d), with the activation of Marijuana itself following closely behind. The second important thing to happen is that the inferences from the second phrase reach Inside-Of-Opaque. This allows Inside-Of-Opaque to start receiving activation from Block-See (Figure 9c) through its results-in branch, causing an immediate leap in Inside-Of-Opaque's incoming activation at cycle 93 (Figure 10b)[10].

As time goes on, feedback through the virtual structure between the Police-Capture frames and their Marijuana filler enables Marijuana to overtake Cooking-Pot (cycle 160), thus lexically reinterpreting the meaning of the word *pot*. Similarly, the new evidence from the Transfer-Self ↔ Proximity-Of ↔ See-Object ↔ Block-See path causes Inside-Of-Opaque to accumulate more evidential activation than Inside-Of-Dishwasher, thus making a pragmatic reinterpretation of John's reason for putting the pot inside the Dishwasher. The final interpretation of **Hiding Pot** is the most highly-activated path of frames in the network, which at stability includes the Inside-Of-Opaque path and the Police-Capture frames, representing the interpretation that John was trying to avoid the detection of his Marijuana from the police by hiding it in an opaque dishwasher.

5.2. The Effect of Selectional Restrictions

ROBIN's network structure that controls the spread of activation based on selectional restrictions is crucial to controlling crosstalk from logically unrelated inferences. As an example, consider the sentence:

After Bill put the omelette on the stove, he put the bowl inside the dishwasher. (**Cook-and-Clean**)

The most likely interpretation of **Cook-and-Clean** is that Bill put the bowl in the dishwasher so that he could eventually clean it after cooking his omelette. In the network, the instance of Inside-Of-Dish-

[10]In addition to propagating signatures representing role-bindings, ROBIN propagates the signature of the frame the bindings originally came from, i.e. the original sentence's phrase. Structure similar to that enforcing selectional restrictions ensures that each frame only receives evidential activation from at most one frame with the same phrase origination, preventing potentially ruinous effects from circular self-feedback loops. Here, for instance, Inside-Of-Opaque initially receives no evidential activation from Block-See, since the sole evidence for Block-See came through Inside-Of-Opaque itself (from **P2**). By cycle 93, however, Block-See has evidence on its own right from **P2** (through See-Object), so its evidential activation can start reinforcing Inside-Of-Opaque.

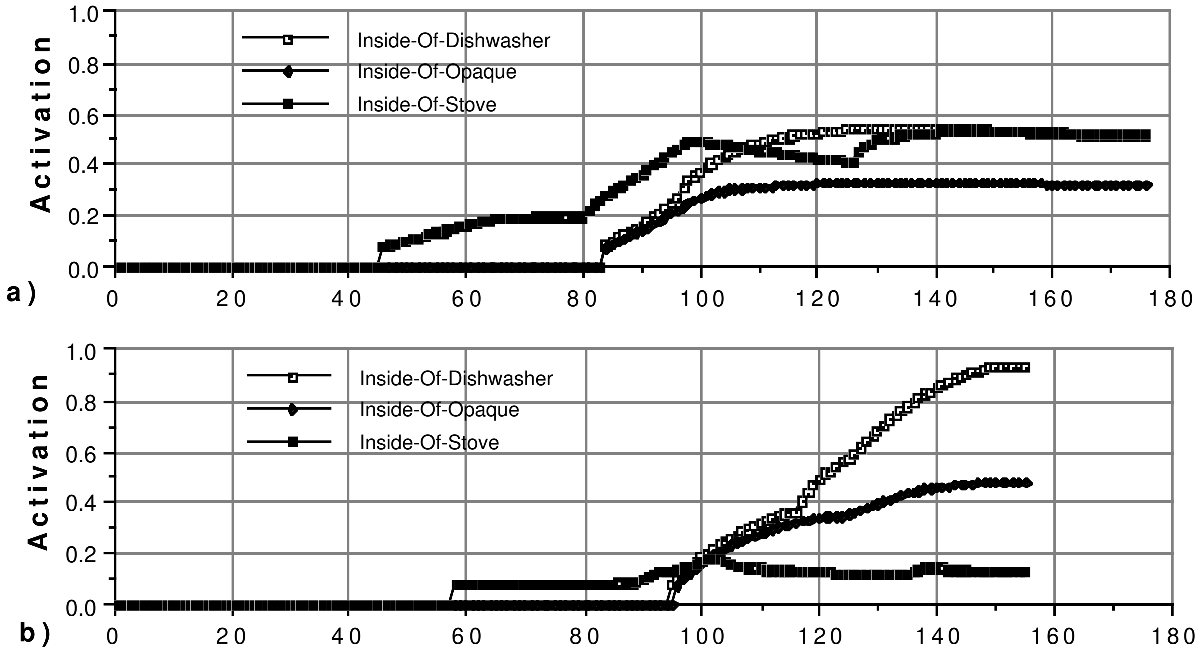

Figure 11. Activations of refinements of Inside-Of after being presented with input for *After Bill put the omelette on the stove, he put the bowl inside the dishwasher.* For **a)** a ROBIN network without structure enforcing selectional restrictions, and **b)** a ROBIN network with selectional restrictions.

washer with Bill as the Actor and Bowl as the Object should thus be the winning refinement of Inside-Of. However, if Inside-Of-Stove is allowed to receive activation from Inside-Of even though its binding constraints are violated, there is a good chance that it could become more activated than Inside-Of-Dishwasher, due to the combined activation from Inside-Of and Stove. As can be seen in Figure 11a, this in fact happens in a network without the structure enforcing selectional restrictions. This "dumb" network has arrived at the decision that Bill was trying to cook something in the bowl when he put it in the dishwasher.

Of course, ROBIN could use the solution of marker-passing systems and evaluate the paths returned by the propagation of signatures. This evaluation would reveal that Inside-Of-Stove's selectional restrictions had been violated (since the Location was actually a Dishwasher). ROBIN could then reject that path and choose the next most-highly activated one, which in this case is the correct Inside-Of-Dishwasher path.

However, one of ROBIN's main goals is to avoid the use of a separate path evaluation mechanism by returning a single, logically-possible inference path. And as can be seen in Figure 11b, this is exactly what happens for **Cook-and-Clean** in a normal ROBIN network with the structure enforcing selectional restrictions. With the selectional restrictions, Inside-Of-Stove still becomes partially activated because of evidence from its prototypical fillers (particularly Stove). But since it does not receive any signatures or evidential activation from Inside-Of, its activation does not come close to competing with Inside-Of-Dishwasher. The network thus settles on the correct interpretation of the sentence.

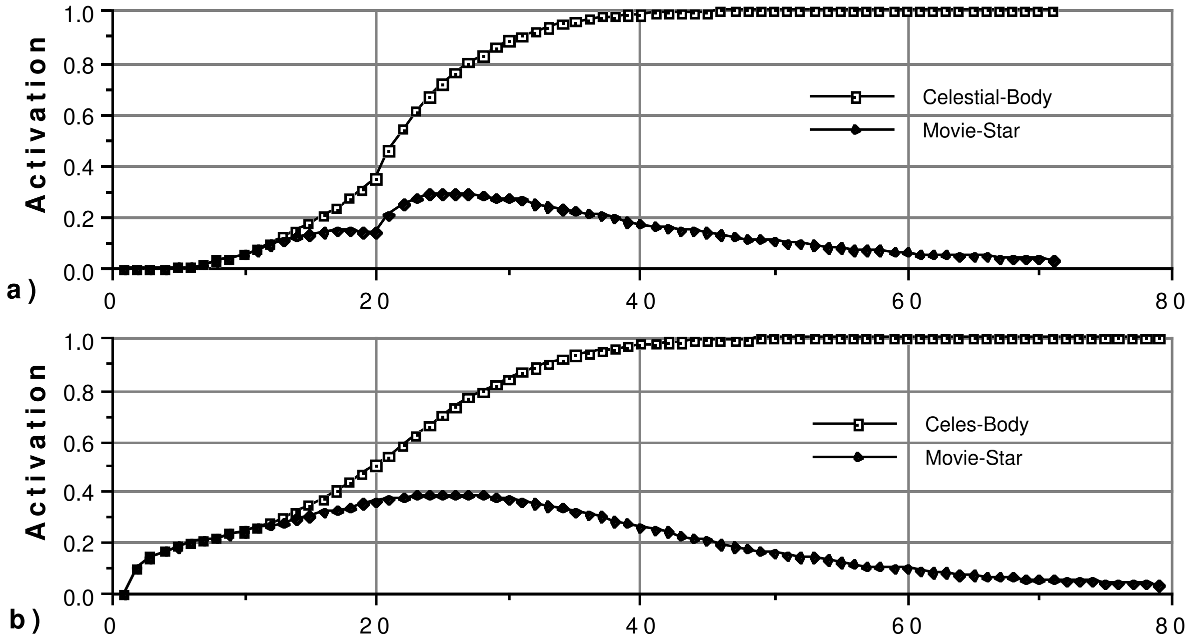

Figure 12. Activations of Celestial-Body and Movie-Star in the network of Figure 2 based on [Waltz & Pollack, 1985] after presentation of input for **a)** *The astronomer saw the star*, and **b)** *The star saw the astronomer.*

5.3. Comparison to Structured Spreading-Activation Models

ROBIN's disambiguation performance is similar to that of Cottrell & Small [1982] and Waltz & Pollack's [1985] structured spreading-activation models on the simple examples that they handle successfully. Of course, signatures allow ROBIN to perform disambiguation on more complex text that requires inferencing, and so is not limited to performing disambiguation based upon the surface semantics of the input.

To compare ROBIN's disambiguation abilities to that of other spreading-activation models and show how important its virtual role-filler structure is for even simple examples that do not require inferencing, consider what happens in the extended Waltz & Pollack network of Figure 2 when it is presented with input for the sentences *The astronomer saw the star* and *The star saw the astronomer.* As shown in Figure 12a, the network disambiguates *star* to mean a Celestial-Body in the sentence *The astronomer saw the star* because of activation flowing over the path Astronomer ↔ Astronomy ↔ Celestial-Body. This is a reasonable disambiguation, since it is the job of astronomers to study celestial bodies. However, when presented with input for *The star saw the astronomer*, the network again disambiguates *star* to Celestial-Body (Figure 12b)[11]. This is quite unfortunate, since celestial bodies lack eyes and so cannot see. The problem is that the network of Figure 2 has no way to recognize the difference between the two sentences, and so provides activation evidence to Celestial-Body through See's Object role, even though it is the Astronomer that is (or should be) bound to it.

[11]It takes slightly longer since *"star"* is clamped first and thus provides early evidence to Movie-Star.

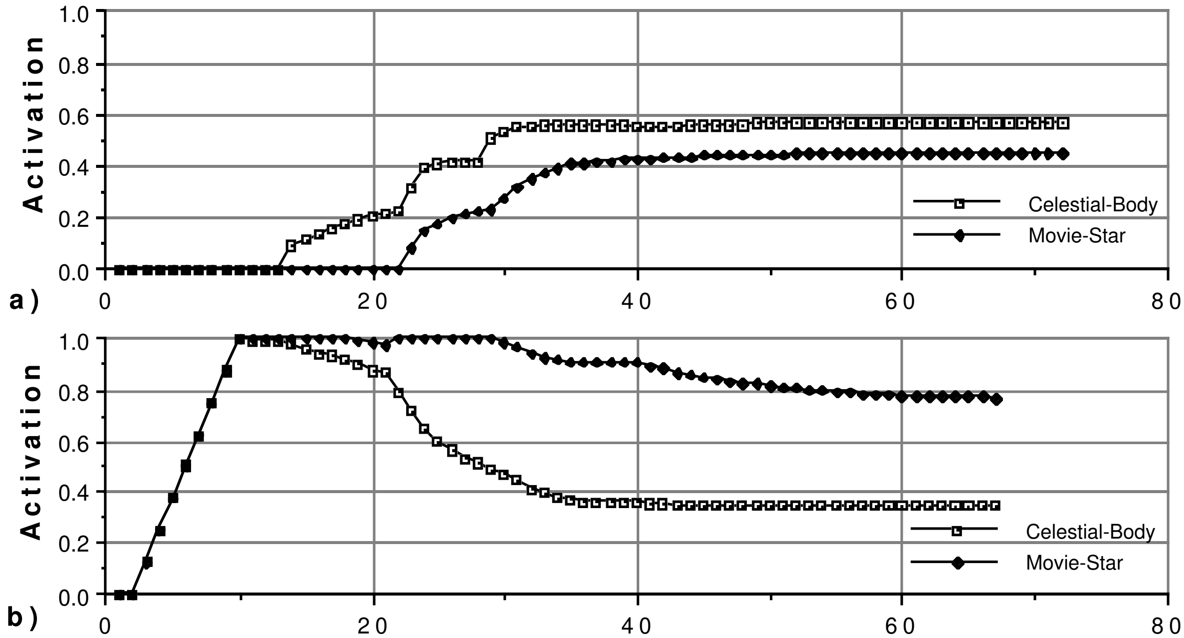

Figure 13. Evidential activations of Celestial-Body and Movie-Star in ROBIN network after presentation of input for **a)** *The astronomer saw the star*, and **b)** *The star saw the astronomer.*

This is a simple example where the combination of ROBIN's signatures, selectional restrictions, and virtual role-filler structure are crucial for successful interpretation. As seen in Figure 13a, ROBIN also disambiguates *star* to Celestial-Body for *The astronomer saw the star.* However, when input for *The star saw the astronomer* is clamped, only the signature of Movie-Star propagates to the Actor role of frame See, since Celestial-Body violates its selectional restrictions. Because of this, only Movie-Star receives activation feedback from See, thus becoming the clear winner and disambiguating the sentence correctly (Figure 13b).

Another difference between ROBIN and other spreading-activation models is the final levels of activation when the network settles. As can be seen in Figures 3 and 12, the direct inhibitory connections of Waltz & Pollack's model drive the non-winner's activations down to zero. This makes it nearly impossible for new context to overcome the inhibition from the winning concept and allow reinterpretation. ROBIN's global inhibition mechanism, on the other hand, serves only to control the spread of activation by normalizing the evidential activations of each concept, leaving their final levels of activation at a value relative to the amount of evidence available for them in that context (e.g. Figure 13). In the case of **Hiding Pot**, this is what allowed the word *pot* to eventually be reinterpreted to Marijuana, despite Cooking-Pot's early dominance. We did not even attempt to run the network with direct inhibitory connections between those competing concepts, because it was clear from the activation plots of Figures 10c and 10d that the new context from *because the police were coming* served only to eventually push Marijuana's net input above that of Cooking-Pot, and would never have been able to overcome the large amount of inhibition that would have been emanating from Cooking-Pot with direct inhibitory connections.

5.4. Comparison to Marker-Passing Models

One way to look at ROBIN's signatures is as a connectionist implementation of a restricted class of symbolic markers, since both signatures and markers allow variable-bindings to be represented and propagated in parallel across semantic networks. Markers, however, can be much more complex than signatures, since they are true symbolic pointers often holding structured information and which can be operated on by symbolic functions on marker-passing nodes. Marker-passing systems also often use multiple types of markers in different stages of the spreading process (e.g. "activation" and "prediction" markers in [Riesbeck & Martin, 1986]) and use named links which may act differently depending upon the type of marker. The converse is that although signatures are less powerful than symbolic markers, they allow ROBIN to perform inferencing using simpler, activation-based connectionist units that implement a single domain and knowledge-independent mechanism.

Propagation of markers and signatures allow both marker-passing systems and ROBIN to generate candidate interpretations of input text in parallel, a crucial advantage over symbolic rule-based models that do so serially. The most important difference between ROBIN and marker-passing systems in terms of disambiguation and reinterpretation is that ROBIN does not need to use a separate path evaluator to select the most plausible interpretation of the many paths generated. As an example of how ROBIN performs the disambiguation and reinterpretation of a marker-passing system, consider how a ROBIN network built from Eiselt's [1987] knowledge base of Figure 1 processes input for the text:

> *Fred proposed to Wilma. Wilma began to cry. Wilma was saddened by the proposal.*
> (**Marriage**)

After activation is spread, both major inference chains between Propose-Marriage and Cry-Tears are instantiated with signatures, just as they are with Eiselt's markers. Of course, each instantiated frame in ROBIN's network also has a level of evidential activation. The levels of activation for the linchpin frames Happy-Event and Sad-Event are shown in Figure 14a. As can be seen, Happy-Event initially is the winning interpretation. However, after the *Wilma was saddened* phrase is input to the network at about cycle 120, Sad-Event gets more evidence and thus becomes the more highly-activated of the two, reinterpreting the text. No resort to separate evaluation metrics or symbolic buffers to store discarded paths is necessary.

Another advantage of ROBIN's connectionist networks over marker-passing networks is that its weighted links and graded activation levels allow the most predictive connections to bias the interpretation more than others. For example, the result shown in Figure 14a was in a network biased to consider marriage proposals as happy events (by having a stronger weight from Propose-Marriage to Happy-Event than to Sad-Event). A more "cynical" network can be modelled simply with modified weights that provide different amounts of evidence, thus leading to a different interpretation without changing the actual knowledge or structure of the network. For example, the result shown in Figure 14b is from a network in which marriage proposals are equally as likely to be sad as happy. On hearing of the proposal, it cannot "decide" between Happy-Event and Sad-Event, so their two activations increase at the same rate. However, when the network is presented with the information that Wilma cried (cycle 60), there is more evidence that she was sad (because of a slight weight bias towards Sad-State from Cry-Tears). The final information that *Wilma was saddened* at cycle 120 only confirms the conclusion. Thus, in ROBIN (as in all structured spreading-activation networks), the same network can return different interpretations to the same input when its weights are biased towards different concepts — as opposed to the binary nature of most marker-passing networks' paths.

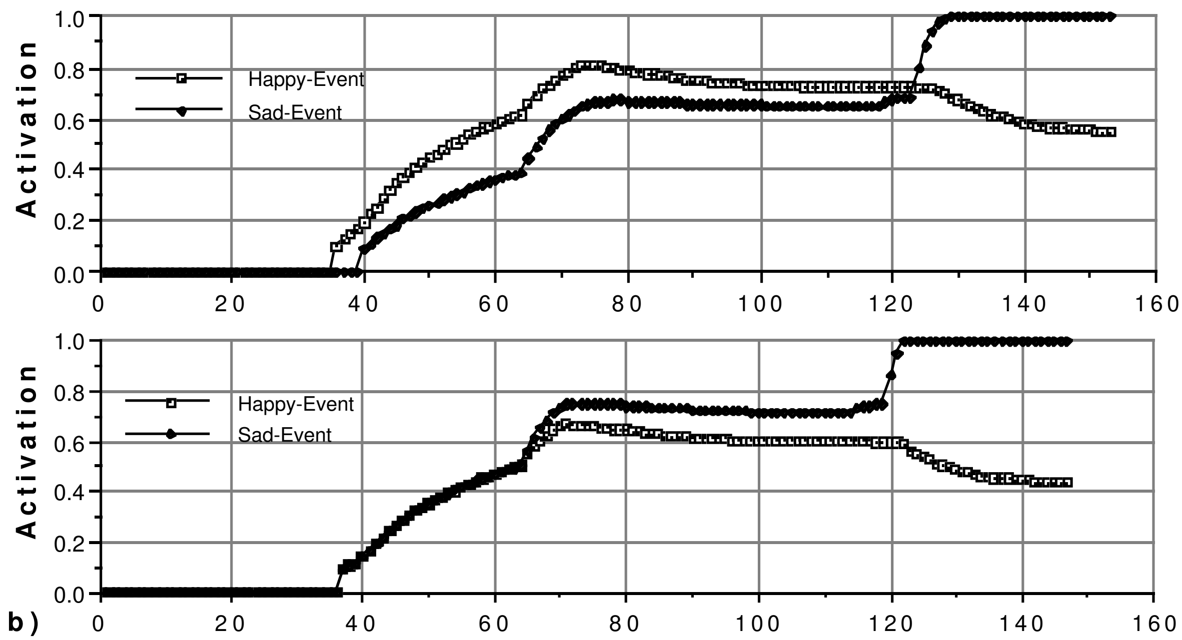

Figure 14. Activations of Happy-State and Sad-State in ROBIN network encoding network of Figure 2 from [Eiselt, 1987], after presentation of input for *Fred asked Wilma to marry him. Wilma began to cry. She was saddened by the proposal.* **a)** network biased towards marriage proposals being happy events (weight from Propose-Marriage to Happy-Event = 0.7 and to Sad-Event = 0.3). **b)** network in which marriage proposals are equally likely to be sad events (weight from Propose-Marriage to Happy-Event and Sad-Event = 0.5).

One interesting thing to note is the gross similarities between the evaluation rules used by marker-passing systems and the kinds of inference paths favored by ROBIN's spread of evidential activation. For example, most marker-passing systems have a rule that selects paths that include more of the input than others. This naturally tends to occur in ROBIN, because each clamped input is a maximal source of evidential activation — so paths that include more of the input will generally have more activation. Other typical marker-passing path evaluation rules and how they seem to correspond to the properties of spreading activation in ROBIN are shown in Table 1. Of course, while marker-passing evaluation heuristics are hard and fast rules, the corresponding tendencies in ROBIN are soft constraints that emerge from the spreading-activation process. These emergent constraints can be and often are "overruled" by other activation tendencies and biases from connection strengths and priming. Most important of all is that ROBIN's disambiguation and reinterpretation happens within the network at the same time as inference paths are generated. In marker-passing systems, on the other hand, path evaluators are a symbolic mechanism separate from the spreading-activation process that operates in serial after paths have been generated, a huge disadvantage as the size of the networks increase and the number of generated inference paths to be evaluated explodes.

5.5. Future Work

In the future, there are four main areas that we would like to explore: (1) the ability to handle embedded role-bindings, (2) increasing the network's capacity, (3) formation of long-term episodic memory, and

Typical Marker Propagation Rules	Propagation of Signatures
Only propagate markers across a certain distance of nodes.	Evidential activation diminishes the further away it gets from an input source. Once it goes below threshold, it and signatures stop propagating.
Typical Path Evaluation Rules	**Evidential Activation Effect**
Reject paths whose frames' binding constraints have been violated.	Selectional restriction structure stops signatures and evidential activation from spreading to frames whose binding constraints would be violated.
Select paths that include most of the input.	Each clamped input is a maximal source of evidential activation — so paths that include more of the input generally have more activation.
For paths explaining the same inputs, select the shortest path.	Shorter paths have less distance for activation to propagate away and diminish, so generally have more activation.

Table 1. Typical marker-passing rules for propagation of markers and evaluation of inference paths contrasted with a similar gross effect from the spread of activation in ROBIN.

(4) the ability to determine the initial surface role-bindings with additional lexical information in the networks.

5.5.1. Embedded Role-Bindings

Using signatures of pre-existing concepts, ROBIN can create and infer novel network instances. However, ROBIN currently cannot dynamically generate and propagate *new* signatures for one of these instances. This ability is crucial for recursive structures, such as in: *John told Bill that Fred told Mary that...* Here each Object of the telling is itself a novel frame instance not having a pre-existing signature. We are currently exploring a solution in which the signatures of the newly-instantiated frames themselves are propagated, a solution that is analogous to that of [Ajjanagadde, 1990] for the phase-clock binding approach. Until a successful solution for embedded signatures is found, ROBIN's inferencing capabilities will be somewhat limited in comparison to symbolic rule-based systems[12].

5.5.2. Network Capacity

ROBIN currently only has the capacity to understand examples of from one to three sentences in length, such as **Hiding Pot** and **Marriage**. A major limitation of the model as described is that each frame can only have one instance at any given time, since binding units can only hold a single signature activation at once. Because of this, ROBIN cannot represent or interpret any stories involving two different seeing or eating events, for instance. We are currently exploring a solution in which each frame will have more than one set of conceptual and binding units, each capable of holding a separate dynamic instantiation.

However, even with the capacity to hold multiple instances of each frame, the network's capacity will still be limited by the fact that stories' interpretations are represented as activation across a finite network. The evidential activation of original parts of a story that are not bolstered by new context will decay away and be lost as time progresses. This is not a problem in marker-passing networks, since they can simply

[12]But not limited in comparison to other structured connectionist models of disambiguation, which cannot handle even simple role-bindings.

store the generated inference paths incrementally in a separate symbolic buffer[13]. Of course, we could use such a solution for ROBIN, but we would prefer to find a purely connectionist solution. Part of the question will be exploring how much of a story can or should be held by the short-term memory of activation in the network before it decays away. A couple of sentences? A couple of paragraphs? On the order of psychological seconds, minutes, or hours?

5.5.3. Formation of Long-Term Episodic Memory

Signatures allow ROBIN to create novel network instances over its pre-existing structure, but the activation of these instances is transient. Over time, repeated instantiations should cause modification of weights and recruitment of underutilized nodes [Diederich, 1990] to alter network structure. Possible methods of storing the short-term memory inferred instances in long-term episodic memory by some kind of distributed learning mechanism must also be explored, likely in conjunction with the use of signatures as distributed patterns of activation, rather than the simplified arbitrary real values used here.

5.5.4. Lexical Information and Initial Role-Bindings

ROBIN does not currently address the problem of deciding upon the original syntactic bindings, i.e. that *pot* is bound to the Object role of the phrase. Rather, ROBIN'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 ROBIN is currently given by hand.

6. CONCLUSIONS

A major source of uncertainty for natural language understanding systems is lexical and pragmatic ambiguity. Symbolic, rule-based models can make high-level inferences necessary for understanding text, but handle ambiguity poorly, especially when later context requires a reinterpretation of the input. Distributed connectionist models, on the other hand, are able to learn to perform disambiguation, but only for simple sentences that can be understood based on the surface semantics of the input or on script-based stories that they have been previously trained to recognize.

Marker-passing networks and structured spreading-activation networks seem to be better suited to disambiguation of text requiring inferencing. Marker-passing systems use their built-in symbolic abilities to perform inferencing and generate possible interpretations of the text in parallel. However, they must employ separate path evaluation mechanisms to decide between the multiple paths generated by propagation of markers. Worse yet, the number of inference paths generated explodes as the size of knowledge-bases increase, slowing them down dramatically. Structured spreading-activation networks, on the other hand, use their weighted connections and graded levels of activation to select a single most-plausible interpretation in a given context through the spreading-activation process. Unfortunately, because of their inability to represent variable bindings and perform inferencing, they have been unable to go beyond disambiguation based on the surface semantics of the input.

We have described a structured spreading-activation model, ROBIN, that is able to perform much of the massively-parallel inferencing of marker-passing systems by propagating activation patterns serving as

[13]Though they still face the problem of determining when to remove individual markers from the network.

concepts' *signatures*. Most importantly, the ambiguous candidate interpretations generated by the propagation of signatures are selected between by the spread of activation along ROBIN's *evidential* semantic network structure, which is similar to that of normal spreading-activation networks. The network thus combines evidence from context for each inference path and settles upon a single most-highly activated path by constraint satisfaction. Furthermore, because each concept in the network retains a level of activation that corresponds to the amount of evidence in its favor, reinterpretation occurs automatically if new context causes another inference path to become more highly-activated than the original winner.

Once structured spreading-activation networks are extended to handle inferencing, potential problems from crosstalk make it vital for their static evidential structure to interact with the dynamic structure represented by the current variable bindings of the network. We have identified two reasons why this is especially important: to assure that activation feedback is between frames and their actual (rather than prototypical) role-fillers, and to stop activation from spreading to frames whose selectional restrictions have been violated. ROBIN assures this *virtual structure* by connections of nodes and links between the signature and evidential portions of the network, thus eliminating large sources of potential crosstalk.

REFERENCES

Ajjanagadde, V. (1990): Reasoning With Function Symbols in a Connectionist System. *Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*. Boston, MA, July 1990.

Ajjanagadde, V. & Shastri, L. (1989): Efficient Inference With Multi-Place Predicates and Variables in a Connectionist System. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society*. Ann Arbor, MI, August 1989.

Barnden, J. & Srinivas, K. (1990): Overcoming Rule-Based Rigidity and Connectionist Limitations Through Massively-Parallel Case-Based Reasoning. *International Journal of Man-Machine Studies*, this issue.

Charniak, E. (1986): A Neat Theory of Marker Passing. *Proceedings of the Fifth National Conference on Artificial Intelligence (AAAI-86)*, Philadelphia, PA, August, 1986.

Cottrell, G. & Small, S. (1982): A Connectionist Scheme for Modeling Word-Sense Disambiguation. *Cognition and Brain Theory, 6*, p. 89-120.

Diederich, J. (1990): Steps Toward Knowledge-Intensive Connectionist Learning. In J. Barnden and J. Pollack (eds.), *Advances in Connectionist and Neural Computation Theory*, Ablex Publishing. In press.

Dyer, M. G. (1983): *In-Depth Understanding: A Computer Model of Integrated Processing for Narrative Comprehension*. Cambridge, MA: The MIT Press.

Eiselt, K. P. (1987): Recovering From Erroneous Inferences. *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, WA, July 1987.

Feldman, J. A. & Ballard, D. H. (1982): Connectionist Models and Their Properties. *Cognitive Science, 6 (3)*, p. 205-254.

Granger, R. H. (1980): When Expectation Fails: Towards a Self-Correcting Inference System. *Proceedings of the First Annual National Conference on Artificial Intelligence*, Stanford, CA.

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*, p. 227-246. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Hendler, J. (1989): Marker-Passing Over Microfeatures: Towards a Hybrid Symbolic/Connectionist Model. *Cognitive Science, 13 (1)*, p. 79-106.

Hirst, G. (1983): *Semantic Interpretation Against Ambiguity.* Ph.D. Dissertation, Brown University.

Holldobler, S. (1990): A Structured Connectionist Unification Algorithm. *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-90)*, Boston, MA, July 1990.

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.

Lange, T. & Dyer, M. G. (1988): Dynamic, Non-Local Role-Bindings and Inferencing in a Localist Network for Natural Language Understanding. In David S. Touretzky (ed.), *Advances in Neural Information Processing Systems I*, p. 545-552. San Mateo, CA: Morgan Kaufmann (Collected papers of the IEEE Conference on Neural Information Processing Systems — Natural and Synthetic, Denver, CO, November 1988).

Lange, T. & Dyer, M. G. (1989): High-Level Inferencing in a Connectionist Network. *Connection Science, 1 (2)*, p. 181-217.

Lange, T., Hodges, J., Fuenmayor, M., & Belyaev, L. (1989): DESCARTES: Development Environment For Simulating Hybrid Connectionist Architectures. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*, Ann Arbor, MI, August 1989.

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. & Dyer, M. G. (1989): A Modular Neural Network Architecture for Sequential Paraphrasing of Script-Based Stories. *Proceedings of the International Joint Conference on Neural Networks*, Washington, DC, June, 1989.

Norvig, P. (1989): Marker Passing as a Weak Method for Text Inferencing. *Cognitive Science, 13 (4),* p. 569-620.

Quillian, M. R. (1969): The Teachable Language Comprehender: A Simulation Program and Theory of Language. *Communications of the ACM, 12*, p. 459-476.

Riesbeck, C. K. & Martin, C. E. (1986): Direct Memory Access Parsing. In J. Kolodner and C. Riesbeck (eds.), *Experience, Memory, and Reasoning*, p. 209-226. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Riesbeck, C. K. (1975): Conceptual Analysis. In R. Schank (ed.), *Conceptual Information Processing*, New York: American Elsevier.

Schank, R. C. & Abelson, R. (1977): *Scripts, Plans, Goals and Understanding*. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

St. John, M. F. (1990): *The Story Gestalt: Text Comprehension by Cue-based Constraint Satisfaction.* Technical Report # 9004, Department of Cognitive Science, University of California, San Diego.

Touretzky, D. S. (1990). Connectionism and Compositional Semantics. In J. A. Barnden and J. B. Pollack (eds.), *Advances in Connectionist and Neural Computation Theory*, Norwood, NJ: Ablex.

Touretzky, D. S. & Hinton, G. E. (1988): A Distributed Connectionist Production System. *Cognitive Science, Vol. 12 (3),* p. 423-466.

Waltz, D. & Pollack, J. (1985): Massively Parallel Parsing: A Strongly Interactive Model of Natural Language Interpretation. *Cognitive Science, 9 (1)*, p. 51-74.