Resolving lexical ambiguity computationally with spreading activation and Polaroid Words

Graeme Hirst
Department of Computer Science
University of Toronto
Toronto, Canada M5S 3G4
gh@cs.toronto.edu

Originally published in: Small, Steven L.; Cottrell, Garrison W.; and Tanenhaus, Michael K. (editors). *Lexical ambiguity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence*. San Mateo, CA: Morgan Kaufmann Publishers, 1988. 73–107.

1 Introduction

Any computer system for understanding natural language input (even in relatively weak senses of the word *understanding*) needs to be able to resolve lexical ambiguities. In this paper, I describe the lexical ambiguity resolution component of one such system.

The basic strategy used for disambiguation is to "do it the way people do". While cognitive modeling is not the primary goal of this work, it is often a good strategy in artificial intelligence to consider cognitive modeling anyway; finding out how people do something and trying to copy them is a good way to get a program to do the same thing. In developing the system below, I was strongly influenced by psycholinguistic research on lexical access and negative priming—in particular by the results of Swinney [1979], Seidenberg, Tanenhaus, Leiman, and Bienkowski [1982], and Reder [1983]. In section 5 below, I will discuss the degree to which the system is a model of ambiguity resolution in people.

The system adheres to the following principles:

- Disambiguation cues come from many sources: syntax, word meaning, context, and knowledge of the world. The resolution component must be able to appeal to any or all of them.

- In general, all possible meanings of a word should be considered. Earlier computational methods such as scripts [Schank and Abelson 1977], in which the context pre-selects allowable meanings for ambiguities, are much too inflexible, and are in conflict with the experimental results mentioned above.

- Resolution happens as soon as possible. If sufficient disambiguating information precedes the word, then resolution occurs immediately; otherwise, it happens as soon as subsequent words have provided enough information, and in any case it must occur by the end of the sentence. It is clear that people work like this; they don't, for example, wait until a sentence is complete and then go back and start resolving the lexical ambiguities.

- Determining which case slot is flagged by a preposition or a particular syntactic position is a process not unlike the disambiguation of content words, and should be handled as far as possible by the same mechanism. That is, both tasks are a kind of lexical disambiguation, so if one process can deal with some or all of both, then, by Occam's Razor, we should prefer it over two independent processes.

Despite our respect for psychological reality as a design strategy (and not as an end in itself), the resolution component is part of an artificial intelligence system, and must therefore also be able to work with the other components of the system; compromises were sometimes necessary.

This work is part of a project on semantics and ambiguity in natural language understanding. Other components of the system include a syntactic parser, named Paragram [Charniak 1983a], based on that of Marcus [1980]; a semantic interpreter named Absity [Hirst 1983, 1987]; a structural disambiguator called the Semantic Enquiry Desk [Hirst 1984a, 1987]; and a knowledge representation and inference system called Frail [Wong 1981a, 1981b; Charniak, Gavin, and Hendler 1983].

Absity is a compositional semantic interpreter that runs in tandem with the parser. Every time the parser creates a new syntactic structure from words or from smaller structures, Absity constructs a semantic object to correspond to that structure, using the semantic objects that correspond to the words or smaller parts. Thus starting from words and their meanings, parsing and semantic interpretation proceed in lockstep, and there is always a well-formed semantic representation for every partial or complete syntactic structure. If the sentence is structurally ambiguous, that is, if the parser has to make a choice between two structures, it will ask the Semantic Enquiry Desk to decide which alternative is best. The semantic objects are represented in the Frail frame language. For the purposes of the present paper, Frail may be thought of as a fairly conventional representation composed of interconnected schemas, with an associated inference and retrieval mechanism; Hirst [1987] discusses in detail the adequacy of such representations as a target for semantic interpretation.

There is an immediate catch in this interpretation scheme. Occasionally a word cannot be disambiguated until some time after its occurrence, whereas Absity wants its semantic object as soon as the word appears. But usually the meaning of a single ambiguous word does not affect the immediate processing of the next few words after it. What we do, therefore, is give Absity a **fake** semantic object, with the promise that in due course it shall be replaced by the real thing. The fake is labeled with everything else that Absity needs to know about the object (such as its syntactic category or possible categories), Absity builds its semantic structure with the fake, and when the real object is available, it is just slipped in where the fake is.

The fakes that we give Absity can be thought of as self-developing Polaroid[1] photographs of the semantic object, and the promise is that by the time the sentence is complete, the photograph will be a fully developed "picture" of the desired semantic object. Even as the picture develops, Absity is able to manipulate the photograph, build it into a structure, and indeed do everything with it that it could do with a fully developed photograph, except look at the final picture. Moreover, like real Polaroid photographs, these have the property that as development takes place, the partly developed picture will be viewable and usable in its intermediate form. That is, just as one can look at a partly developed Polaroid picture and determine whether it is a picture of a person or a mountain range, but perhaps not which person or which mountain range, so it is possible to look at a **Polaroid Word** and get an idea of what the semantic object it shows looks like.

I will describe the operation of Polaroid Words in section 3. Before that, in section 2, I discuss marker passing, a mechanism that Polaroid Words uses for finding associations between words.[2] In section 4, I discuss some of the ways in which Polaroid Words are not yet adequate, and then, in section 5, look at the extent to which Polaroid Words are a psychological model. I assume throughout this chapter that the input sentence is not structurally ambiguous; in Hirst [1987] I show how Polaroid Words work together with the Semantic Enquiry Desk, each constraining the other, when the correct parse of the sentence cannot be determined immediately.

2 Marker passing

It is well known that an association between one sense of an ambiguous word and other words in the sentence or context can be an important disambiguation cue. Psycholinguistic research on lexical disambiguation has shown that semantic priming—that is, the previous occurrence of an associated word—speeds up people's

[1] *Polaroid* is a trademark of the Polaroid Corporation for its range of self-developing photographic products and other products. The word is used here to emphasize the metaphor of a self-developing semantic object, and the system described herein carries no association with, or endorsement by, the Polaroid Corporation.

[2] These sections differ in some details from the description of an earlier design in Hirst and Charniak [1982].

disambiguation, and may lead the retrieval process straight to the correct meaning.[3]

The lexical disambiguation program for Absity therefore needs a mechanism that will allow it to find semantic associations. Quillian suggested as long ago as 1962 that following connections in a semantic network is an excellent way of finding such associations. Our mechanism for this will be **marker passing** in the Frail knowledge base.

Marker passing can be thought of as passing tags or markers along the arcs of the knowledge base, from frame to frame, from slot to filler, under the rules to be discussed below. It is a discrete computational analogue of **spreading activation** in semantic memory. The network of frames or schemas corresponds to the conceptual and lexical network of semantic memory, and a connection implies a semantic relationship of some kind between its two nodes. Passing a marker from one node to another corresponds to activating the receiving node. If marker passing is breadth-first from the starting point (new markers being created if a node wishes to pass to two or more other nodes simultaneously), then marker passing will "spread" much as spreading activation does.[4]

Marker passing was first used in artificial intelligence by Quillian [1968, 1969], who used it to find connections between concepts in a semantic network. Marker passing is, of course, expensive when the net is interestingly large. Fahlman [1979], who used it for deduction in his NETL system, proposed super-parallel hardware for marker passing. Although the present scheme is much simpler than Fahlman's, I too assume that hardware of the future will, like people of the present, be able to derive connections between concepts in parallel, and that the serial implementation to be described below is only an interim measure.

2.1 Marker passing in Frail

The frame language Frail contains a built-in marker passer (*MP* for short) that operates upon the Frail knowledge base [Charniak, Gavin, and Hendler 1983]. The MP is called with the name of a node (a frame, slot, or instance) as its argument to use as a starting point. From this origin, it marks all nodes in the knowledge

[3] Or perhaps straight to an incorrect meaning, if the semantic prime is misleading; see section 5.1.

[4] Note that this is a very localist system in which each node represents a concept. It is thus in contrast to distributed connectionist-style systems, in which a single concept is represented by a **collection** of nodes among which activation may pass.

base that participate in assertions that also contain the origin; these can include slots, slot restrictions, and IS-A relationships. For example, suppose the origin is to be the frame that describes libraries:

(1)
```
[frame:  library
   isa:  institution
 slots:  (function (store-books lend-books))
         (employee (librarian))
         ... ]
```

Markers would be placed on `institution`, `store-books`, `lend-books`, `employee`, `librarian`, and so on. Markers take the form of a list of node names interleaved with the connection that permitted the marker to be placed, so that there is always information to retrace the path.

Once all the nodes reachable in one step from the origin are marked, each node reachable from these nodes—that is, each node two steps from the origin—is marked. For example, if a slot in the frame `store-books` contains `book` (as it surely would), then `book` would be marked. Thus marker passing proceeds, fanning out from the origin until all nodes whose distance is *n* or less from the origin have been marked, where *n* defaults to 5 if the programmer doesn't specify otherwise.[5]

If at any time during marker passing the MP comes upon a node already marked by a previous call, then a **path** (or **chain**) has been found between the origin node of the present call and that of a previous call. The MP uses the pre-existing mark on the node to construct the full origin-to-origin path. Suppose that in the example above, the `book` node had been found to be already marked, indicating a connection to it from `publisher`; the MP can then note that a path has been found between `library` and `publisher`. It is also possible that the origin itself has been marked by a previous call to the MP, resulting in an instantly discovered path. We call such paths **immediate paths** to distinguish them from **constructed paths**, such as that of the example above in which the intersection occurs at a third node.

When marking is finished, the MP returns a list of all the paths (if any) that it found. One may, at any time, clean the markers from all nodes in the knowledge base.

[5] See my remarks in section 5.3 about 'magic numbers' in marker passing and 'spending' activation.

2.2 Lexical disambiguation with marker passing

In this section, I give a very simple example of lexical disambiguation with the Frail marker passer that was discussed in the previous section. In later sections, I will refine and extend these disambiguation mechanisms considerably.

The marker passer operates independently of Absity and in parallel with it. That is, following only morphological analysis, the input sentence goes to both the Paragram parser and the MP, both of which separately grind away on each word as it comes in. Suppose the input is (2), an example chosen especially because it contains several ambiguous words that can be resolved purely by association cues:

(2) Nadia's <u>plane</u> <u>taxied</u> to the <u>terminal</u>.

The words *plane*, *taxi*, and *terminal* are all ambiguous. Note that the ambiguity of *taxi* is categorial: it can be a noun meaning **vehicle with driver for hire**, or a verb meaning (of an airplane) **to travel at low speed on the ground**. Since the MP has no access to syntactic information, it looks at all meanings for each word, regardless of part of speech; marker chains from origins that later prove to be syntactically inappropriate will simply be ignored by other processes.

As the words come in from left to right, the MP passes markers from the frames representing each known meaning of each open-class word in the sentence (including unambiguous ones such as *Nadia*). In (2), short paths would be found between the frames `airplane` and `airport-building`, which were starting points for *plane* and *terminal*, because the latter would be a frame mentioned in one of the activities described in the frame for the former. Likewise, `airplane` and `aircraft-ground-travel` (*plane* and *taxi*) will be connected, because the latter is an activity attributed to the former. These connections indicate that the corresponding meanings of *plane*, *terminal*, and *taxi* should be chosen. (A path will also be found between `airport-building` and `aircraft-ground-travel`, but this gives no new information.) Markers will also be passed from the frames representing the other meanings of *plane*, *taxi*, and *terminal*, namely `wood-smoother`, `taxicab`, and `computer-terminal`, but these paths will go off into the wilderness and never connect with any of the other paths.

2.3 Constraining marker passing

Since marker passing is a blind and mindless process, it is clear that many paths in the knowledge base will be marked besides the ones that provide useful disambiguating information. In fact, if the MP gets too carried away, it will eventually mark everything in the knowledge base, as every node in the base can be reached from any other, and one will then find paths between the wrong senses of ambiguous words as well as between the right senses. For example, a connection could be found between `airplane` and `computer-terminal` simply by passing markers up the IS-A chain from `airplane` through `vehicle` and the like to `mechanical-object`, and then down another IS-A chain from there to `computer-terminal`. Therefore, to prevent as many "uninteresting" and misleading paths as possible, we put certain constraints on the MP and prohibit it from taking certain steps.

First, as I mentioned in section 2.1, Frail passes markers a maximum of n arcs from the origin. One would normally choose n to be small compared to the size of the knowledge base. Second, Frail permits the programmer to specify restrictions on passing markers along various types of path. For example, by default the MP will pass markers only upwards along IS-A links, not downwards—that is, markers are passed to more general concepts, but never to more particular ones (thereby prohibiting the path from `mechanical-object` to `computer-terminal` mentioned above). These restrictions are specified in the form of a predicate supplied by the programmer and attached to the name of the arc. Before attempting to pass a marker, the MP will evaluate the predicate, which has access to the origin, the present node, and the path between them; if the value of the predicate is `nil`, no marker is passed.

Determining exactly what restrictions should be placed on marker passing is a matter for experiment [see Hendler 1986, 1987]. I postulate restrictions such as an **anti-promiscuity rule**: not allowing paths to propagate from nodes with more than c connections, for some chosen c. This is because nodes with many connections tend to be uninteresting ones near the top of the IS-A hierarchy—`mechanical-object`, for example. We must be careful, however, not to be so restrictive that we also prevent the useful paths that we are looking for from occurring. And no matter how clever we are at blocking misleading paths, we must be prepared for the fact that they will occasionally turn up. The problem of such **false positives** is discussed by

Charniak [1983b, 1986], who posits a **path checker** that would filter out many paths that are uninteresting or silly.

3 Polaroid Words

In section 1, I introduced the idea of the Polaroid Word mechanism (*PW* to its friends), which would be responsible for disambiguating each word. As I noted, there are many sources of information that can be used in disambiguation; it is up to the mechanism of the PW to use whatever information is available to it to make a decision for each word. Often, as in the case of example (2) of section 2.2, all that is required is looking at the paths found by the marker passer. At other times, MP will return nothing overwhelmingly conclusive; or, in the case of a polysemous word—one whose several meanings are related—more than one meaning may be marked. It is then necessary for PWs to use other information and negotiate between possible meanings. In this section I will describe in detail this aspect of the operation of Polaroid Words.

3.1 What Polaroid Words look like

While it would be quite possible to operate Polaroid Words under the control of a single supervisory procedure that took the responsibility for the development of each "photograph", it seems more natural instead, because of the parallelism, to put the disambiguation mechanism (and the responsibility) into each individual Polaroid Word. That is, a PW will be a procedure, running in parallel with other PWs,[6] whose job it is to disambiguate a single instance of a word. At this point, however, we find we must stretch the Polaroid photograph metaphor, for unlike that of a real self-developing photograph, a PW's development cannot be completely self-contained; the PWs will have to communicate with one another and with their environment in order to get the information necessary for their disambiguation. So disambiguation is, in effect, a relaxation process. The idea of communicating procedures, one per word, brings to mind Small's word experts [Adriaens and Small 1988]. The similarity between PWs and Small's procedures is, however, only superficial; the differences will become apparent as we describe PWs in detail.

[6] In the implementation described below, only one PW is active at a time, in order to simplify the programming.

```
[slug (noun):
     gastropod-without-shell
     bullet
     metal-stamping
     shot-of-liquor]
```

Figure 1: Packet of knowledge for *slug* for noun Polaroid Word; each line is the name of a frame in the Frail knowledge base.

Instead of having a completely different PW for each word, we have but one kind of PW for each syntactic category; for example, there is a noun PW and a verb PW. Each noun uses the same disambiguation procedure as all the other nouns; each verb uses the same procedure as the other verbs, and similarly for other syntactic categories.[7] The knowledge about the meaning of each individual word is kept distinct from the disambiguation procedure itself, and indeed much of the knowledge used by PWs is obtained from the Frail knowledge base when it is necessary. When a new PW is needed, an instance of the appropriate type is cloned and is given a little packet of knowledge about the word for which it will be responsible. (Sometimes I will be sloppy and call these packets Polaroid Words as well. No confusion should result.) As far as possible, the packets contain only lexical knowledge—that is, only knowledge about how the word is used, rather than world knowledge (already available through Frail) about the properties of the word's denotations.

The simplest packet of knowledge is that for a noun: it just contains a list of the semantic objects—the frames or schemas in the Frail knowledge base—that the noun could represent. Figure 1 shows the knowledge packet for the noun *slug*. Any information needed about properties of the senses of the noun is obtained from the Frail knowledge base.

The packet for prepositions and pseudo-prepositions[8] is a little more complicated; listed with each possi-

[7] At present, PWs are implemented only for nouns, verbs, prepositions, and, in rudimentary form, noun modifiers. Determiners are straightforward, and PWs for them may exist later; see section 3.5.

[8] **Pseudo-prepositions** are the case flags that occur as a syntactic position: *SUBJ*, *OBJ*, and *INDOBJ* for subject, object, and

```
[with (prep):
      instrument  (and physobj (not animate))
      manner      manner-quality
      accompanier physobj]
```

Figure 2: Packet of knowledge for *with* for preposition Polaroid Word; each line is the name of a slot and a corresponding slot restriction predicate.

ble semantic object, whose semantic type is *frame slot*, is a **slot restriction predicate** for each—a predicate that specifies what is required of an instance to be allowed to fill the slot. Figure 2 shows the packet for the preposition *with*; it assumes that the preposition is a case flag. (PWs for prepositions of noun-modifying PPs are discussed in Hirst [1987, section 7.2].) A simple predicate, such as `physobj` ("physical object"), requires that the slot-filler be under the specified node, in this case `physobj`, in the IS-A hierarchy. A complex predicate may specify a boolean combination of features that the filler must satisfy; thus in figure 2, the filler of `instrument` must be a `physobj`, but not an `animate` one.

The predicates listed in the packet for each slot are, in effect, the most restrictive predicate compatible with the restrictions on all instances of the slot for all verbs. In English, for example, an animate entity can never be the INSTRUMENT of an action. Ideally, there would be a process that would automatically compile the preposition information packets from the knowledge base and would help ensure they remain consistent with one another when words are added or changed.

Verbs have the most complex knowledge packets. Figure 3 shows the packet for *operate*. For each meaning, the case slots that it takes are listed, with the preposition or prepositions that may flag each slot. Slot restriction predicates for each slot need not be specified in verb packets, because they may be immediately found in the corresponding frames in the knowledge base. These predicates will, in general, be more restric-

indirect object positions. Pseudo-prepositions are inserted by the parser, and are thereafter treated like other prepositions. See Hirst [1987, section 3.4] for discussion of why this is a good idea.

```
[operate (verb):
  [cause-to-function
        agent       SUBJ
        patient     SUBJ, OBJ
        instrument  SUBJ, with
        method      by
        manner      with
        accompanier with]
  [perform-surgery
        agent       SUBJ
        patient     upon, on
        instrument  with
        method      by
        manner      with
        accompanier with] ]
```

Figure 3: Packet of knowledge for *operate* for verb Polaroid Word; each entry is a frame with a set of slots and the prepositions that may flag each slot.

tive than the predicates given in the PW for the corresponding preposition, but they must, of course, be compatible. For example, in the `perform-surgery` frame, the predicate on `instrument` may be `(property sharp)`, which particularizes the predicate shown for `instrument` in figure 2; a predicate such as `hanim` ("higher animate being") would contradict that in figure 2 and would indicate trouble somewhere. It should be clear that if the semantics are properly characterized, contradictions will not occur, but, again, an automatic system for maintaining consistency would be helpful.

Unlike the other PW knowledge packets, the verb packets contain information that might also be readily obtained from Frail's knowledge base, namely the slots that each verb frame has. Because the knowledge packet has to include a listing of the prepositions that flag each of the verb's slots, the slots themselves have to be listed, necessarily adding a little world knowledge to the word knowledge. The alternative, listing the flags in the Frail definition of the frame in the knowledge base, would just be the same sin at a different site. It might be more elegant if we were able to remove this information from the individual verbs altogether and rather store generalizations about case flags as they relate to the semantic properties of verbs. That is, since verbs are classified in Frail's IS-A hierarchy under such generalizations as `transitive-action` and `transfer-action` in order to support the obvious needs of inference, and since this also provides a nice generalization of slots—for example, all `transfer-actions` have `source` and `destination` slots—we could store a small table that mapped each general verb frame category to a set of flags for its slots. Alas, English is just not quite regular enough to permit this; verbs can get quite idiosyncratic about their case flags. We have already seen in Figure 3 that the two senses of *operate* have different sets of case flags, although both are `transitive-actions`. Another example is the common senses of the verbs *buy* and *sell*, which are often analyzed as referring to the same frame, varying only in how the case flags are mapped to its slots; see Figure 4 for an example. We should not complain about the idiosyncrasy of case flags, however, for it is often a great help in verb disambiguation, especially if the verb is polysemous.

```
[buy (verb):
  [purchase
        destination SUBJ
        source      from
        sold-item   OBJ
        exchange    for
        beneficiary for, INDOBJ] ]

[sell (verb):
  [purchase
        destination to, INDOBJ
        source      SUBJ
        sold-item   OBJ
        exchange    for
        beneficiary for] ]
```

Ross sold the lemming to Nadia.

```
(a ?x (purchase ?x (source=Ross)
                   (destination=Nadia)
                   (sold-item=lemming26)))
```

Nadia bought the lemming from Ross.

```
(a ?x (purchase ?x (source=Ross)
                   (destination=Nadia)
                   (sold-item=lemming26)))
```


Figure 4: Abbreviated packets of knowledge for *buy* and *sell*, using the same frame but different mappings of case flags to slots, and examples of their use in Frail.

3.2 How Polaroid Words operate

PWs operate in parallel with Absity and the parser. Each word comes in to the parser and its syntactic category or possible categories are assigned from the lexicon. A PW process is created for the word for each of its possible syntactic categories; those for categories that prove to be incorrect will not survive. The way the processes work is described below.

There are two easy cases. The first, obviously, is that the word is unambiguous. If this is the case, the PW process merely announces the meaning and uninterestingly hibernates—as soon as PWs have narrowed their possible meanings to just one, they always announce the fact and knock off work. The second easy case is that the marker passer, which has been looking for paths between senses of the new word and unrejected senses of those already seen, finds a nice connection that permits one alternative to be chosen. This was the case with example (2) of section 2.2. I will discuss in section 5.3 exactly what makes a marker passing path "nice". In general, short constructed paths are nice, and immediate paths are nicer.

If neither of these cases obtains, then the PW has to find out some more about the context in which its word occurs and see which of its alternatives fits best. To do this, it looks at certain preceding PWs to see if they can provide disambiguating information; I will describe this process in a moment. Using the information gathered, the PW will eliminate as many of its alternatives as possible. If this leaves just one possibility, it will announce this fact and stop work; if still undecided, it will announce the remaining possibilities, and then sleep until a new word, possibly the bearer of helpful information, comes along.

Communication between PWs is heavily restricted. The only information that a PW may ask of another is what its remaining possibilities are; that is, each may see other partly or fully developed photographs. In addition, a PW is restricted in two ways as to the other PWs it is allowed to communicate with. First, since a sentence is processed from left to right, when a PW is initially invoked it will be the rightmost word in the sentence so far and may only look to PWs on its left. As new words come in, the PW will be able to see them, subject to the second constraint, namely that each PW may look only at its **friends**.[9] Friendships among PWs are defined as follows: Verbs are friends with the prepositions and nouns they dominate; prepositions

[9] Note that friendship constraints do not apply to the marker passer.

are friends with the nouns of their prepositional phrase and with other prepositions; and noun modifiers are friends with the noun they modify. In addition, if a prepositional phrase is a candidate for attachment to a noun phrase, then the preposition is a friend of the head noun of the NP to which it may be attached. The intent of the friendship constraints is to restrict the amount of searching for information that a PW has to do; the constraints reflect the intuition that a word has only a very limited sphere of influence with regard to selectional restrictions and the like, so PW communication is limited to what seems the minimum necessary for the task.

An "announcement" of its meaning possibilities by a PW takes the form of a list of the one or more alternatives from its knowledge packet (with their slot restriction predicates and so on if they are included in the packet) that the PW has not yet eliminated. An announcement is made by posting a notice in an area that all PWs can read; when a PW asks another for its possibilities, what it is actually doing is reading this notice. (PWs only read their friends' notices, of course.)

From the information that the notices provide, a PW eliminates any of its meanings that don't suit its friends. For example, each case slot may occur at most once in a sentence, so if one preposition PW has already decided that it flags, say, the AGENT slot, then a new preposition PW could cross AGENT off its own list. A preposition PW will also eliminate from its list any cases that its dominating verb does not allow it to flag, and any whose predicates are incompatible with its noun complement. Its friends may still be only partly developed, of course, in which case the number of eliminations it can make may be limited. However, if, say, one of its cases requires a `hanim` filler but none of the alternatives in the partly developed noun is `hanim`, then it can confidently cross that case off its list. The PW may use Frail to determine whether a particular sense has the required properties. What is happening here is, of course, very like the use of selectional restriction cues for disambiguation; see Hirst [1987, section 5.5] for discussion of the differences.

Similarly, nouns and verbs can strike from their lists anything that doesn't fit their prepositional friends, and nouns and noun modifiers can make themselves compatible with each other by ensuring that the sense selected for the noun is a frame with which the modifier's sense will fit. (If a PW finds that this leaves it with no alternatives at all, then it is in trouble; this is discussed in section 4.)

When a PW has done all it can for the time being, it announces the result, a fully or partly developed picture, and goes to sleep. The announcement wakes up any of its friends that have not yet made their final decision, and each sees whether the new information—both the new word's announcement and any marker-passing chain between the old word and the new—helps it make up its mind. If so, it too makes an announcement of its new possibilities list, in turn awakening its own friends (which will include the first PW again, if it is as yet undecided). This continues until none can do any more and quiescence is achieved. Then the next word in the sentence comes in, its PW is created, and the sequence is repeated.

3.3 An example of Polaroid Words in action

Let's consider the following example, concentrating on the subordinate clause shown in (4):

(3) Ross found that the slug would operate the vending machine.

(4) SUBJ the slug operate OBJ the vending machine.

Notice that the parser has inserted the pseudo-prepositions *SUBJ* and *OBJ*. We want to work out that *the slug* is a metal stamping, not a gastropod, a bullet, or a shot of whiskey; that the frame that *operate* refers to is `cause-to-function`, not `perform-surgery`; and that *SUBJ* and *OBJ* indicate the slots `instrument` and `patient` respectively. *Vending machine*, we will say, is unambiguous. For simplicity, we will ignore the tense and modality of the verb. The PWs for *slug*, *with*, and *operate* were shown in Figures 1, 2, and 3; those for the other words are shown in Figure 5.

Disambiguation of the subordinate clause proceeds as follows. The first words are *SUBJ* and *slug*; their PWs, when created, have not yet enough information to do anything interesting, nor has marker passing from the senses of *slug* produced anything (since there are no other words with which a connection might be found yet). Then *operate* comes along, and tells the others that it could mean either `cause-to-function` or `perform-surgery`. It too has no way yet of deciding upon its meaning. However, the *SUBJ* PW notices that neither meaning of *operate* uses *SUBJ* to flag a `source` or `destination` case, so it can cross these off its list. It also sees that while both meanings can flag their `agent` with *SUBJ*, both require that the `agent` be

```
[SUBJ (prep):
     agent       animate
     patient     thing
     instrument  physobj
     source      physobj
     destination physobj]

[OBJ (prep):
     patient     thing
     transferee  physobj]

[vending machine (noun):
     vending-machine]
```

Figure 5: Packets of knowledge for *SUBJ*, *OBJ*, and *vending machine*

hanim. None of the possibilities for *slug* has this property, so the *SUBJ* PW can also cross agent off its list, and announce that it means either instrument or patient.

This wakes up the *operate* PW, which notices that only one of its meanings, cause-to-function, can take either an instrument or a patient flagged by *SUBJ*, so it too announces its meaning. The *slug* PW is also woken up, but it is unable to use any of this information.

Next comes the word *OBJ*. It could be patient or transferee, but the verb *operate* doesn't permit the latter, so it announces the former. Note that if *operate* had not already been disambiguated from previous information, this would happen now, as the *operate* PW would notice that only one of its senses takes any case flagged by *OBJ*. Upon hearing that *OBJ* is going to be patient, the PW for *SUBJ* now crosses patient from its own list, since a case slot can appear but once in a sentence; this leaves it with instrument as its meaning. The PW for *slug* is not a friend of that for *OBJ*, so *OBJ*'s announcement does not awaken it. (It is awoken by *SUBJ*'s move, but cannot use the information.)

The noun phrase *vending machine* now arrives, and we assume that it is recognized as a canned phrase representing a single concept [Becker 1975, Wilensky and Arens 1980a, 1980b]. It brings with it a marker-passing chain that, depending on the exact organization of the frames, might be this:

(5) vending-machine $\rightarrow$ coin $\rightarrow$ metal-stamping

since a fact on the vending-machine frame would be that they use coins, and a coin IS-A metal-stamping. This is enough for the *slug* PW to favor metal-stamping as its meaning, and all words are now disambiguated. Now that processing is complete, all markers in the knowledge base are cleared away.

3.4 Recovery from doubt

Now let's consider this example, in which marker passing is not used at all:

(6) The crook operated a pizza parlor.[10]

[10] This is exactly the same meaning of *operate* as in the previous example: cause-to-function. In a context like this, the action is continuous, a matter I ignore here.

This proceeds as example (3) of the previous section did, until *operate* arrives. Since *crook* can either be something that is `hanim`, namely a `criminal`, or not, namely a `shepherd's-staff`, *SUBJ* is unable to make the move that in the previous example disambiguated both it and *operate*, though it can cross `patient` off its list. Still, when *OBJ* comes along, the *operate* PW can immediately eliminate `perform-surgery`. Let us assume that *pizza parlor* is an unambiguous canned phrase, as *vending machine* was. However, after it is processed, the PWs reach a standstill with *SUBJ* and *crook* still undisambiguated, as MP, unaware of current trends in organized crime, finds no connection between *crook* and *pizza parlor*.

If it happens that at the end of the sentence one or more words are not fully disambiguated, then there are three ways that they may yet be resolved. The first is to invoke knowledge of **preferred** or **deprecated meanings** for them—that is, meanings that are especially common or rare. Preferred and deprecated meanings are indicated as an additional part of the knowledge packet for each word; a word can have zero or more of each. For example, the meaning `female-swan` of *pen* is deprecated, and should never be chosen unless there is positive evidence for it (such as strong or weak marker-passing chains—see next paragraph); the meaning `writing-instrument` is preferred, and the meaning `enclosure` is neither preferred nor deprecated. The possibilities that remain are ranked accordingly, and the top one or ones are chosen. In the present example, therefore, the two unfinished PWs look for their preferred meanings. It is clear that in English AGENT is far more common for *SUBJ* than the other remaining possibility, INSTRUMENT, and so the *SUBJ* PW should prefer that. This, in turn will wake up the *crook* PW, which now finds the requirement that its meaning fit *operate*'s `agent`, and therefore chooses `criminal`, completing disambiguation of the sentence.[11]

The second possibility at the end of the sentence is to use "weak" marker-passing chains. It may be the case that during processing of the sentence, marker passing found a path that was considered too weak to be conclusive evidence for a choice. However, now that all the evidence available has been examined and no conclusion has been reached, the weak path is taken as being better than nothing. In particular, a weak path that runs to a deprecated meaning is used as evidence in support of that meaning. In the present implemen-

[11]It is possible that the results will vary depending on which PW applies its preferred meaning first. It is unlikely that there is a single "correct" order for such situations. If a sentence is really so delicately balanced, people probably interpret it as randomly as Polaroid Words do.

tation, the trade-off between weak chains and preferred meanings is accomplished by “magic numbers” (see section 5.3).

If neither preferred meanings nor weak chains help to resolve all the remaining ambiguities, then inference and discourse pragmatics may be invoked. It should be clear that Polaroid Words with marker passing are not a replacement for inference and pragmatics in word sense and case disambiguation; rather, they serve to reduce substantially the number of times that these must be employed. However, there will still be cases where inference must be used. For example, the following sentences couldn’t be disambiguated without inference about the relative aesthetics of factories and flora:

(7) The view from the window would be ruined by the addition of a <u>plant</u> out there.

(8) The view from the window would be improved by the addition of a <u>plant</u> out there.

Similarly, when a president tells us (9):

(9) I am not a <u>crook</u>.[12]

neither marker passing nor Polaroid Words will help us discover that he or she is not denying being a shepherd’s staff, even though we may readily determine that shepherd’s staff he or she is not.[13]

Throughout this process, however, it should be kept in mind that some sentences are genuinely ambiguous to people, and it is therefore inappropriate to take extraordinary measures to resolve residual problems. If reasonable efforts fail, PWs can always ask the user what he or she really meant:

(10) **User:** I need some information on getting rid of moles.

System: What exactly is it that you wish to get rid of? Unsightly skin blemishes, some cute but destructive insectivorous garden pests, uterine debris, unwanted breakwaters, chili and chocolate sauces, or enemy secret agents that have penetrated deep into your organization?

(11) **User:** Are there any planes in stock?

System: We’ve got two carpentry planes and a Boeing 747.

[12]Nixon, Richard Milhous. 11 November 1973.

[13]The present implementation does not have such an inference or pragmatics system available to it.

(The system does not actually have such a natural language response component.)

One particular case of genuine ambiguity corresponds to when PWs **deadlock**. Deadlock between two (or more) PWs is possible if one says "I can be X if you are A, and I can be Y if you are B", while the other says, conversely, "I can be A if you are X, and I can be B if you are Y". In other words, the sentence has two readings, corresponding to the choices $X+A$ and $Y+B$. This can happen if two "parallel" MP paths are found:

(12) Ross was escorted from the <u>bar</u> to the <u>dock</u>. *(a courtroom scene or a harborside scene)*

(13) Each <u>bill</u> requires a <u>check</u>.[14] *(each* **invoice** *requires a* **negotiable instrument***, or each* **proposed law** *requires a* **verification of correctness***, or various other combinations)*

Deadlock cases are probably very rare—it is hard to construct good examples even out of context, let alone in a context—and PWs have no special mechanism for dealing with them.

3.5 Cues unused

Many words can be disambiguated by syntactic cues; for example, many of the various senses of *keep* may be distinguished by the syntactic form of the object.[15] At present, Polaroid Words do not have this sensitivity, because of limited communication between PWs and the Paragram parser. I do not anticipate major difficulties in changing this in future versions of Polaroid Words.

Because PWs are not sensitive to syntax yet, they cannot yet handle passive sentences. Also awaiting this sensitivity are PWs for determiners. For example, the word *the* translates as either `the` or `the-pl`, depending on whether its NP is marked as singular or plural. Determiner PWs would be particularly simple, as they do not have to deal with marker passing nor provide feedback to any other PWs. A sensitivity to syntax would also assist the resolution of nouns such as *rock*, some senses of which can be pluralized and others of which cannot.

[14] Readers in countries where the spelling *cheque* is used should pretend that the example was read aloud to them.

[15] For example, the word *keep* may mean **maintain** if followed by a direct object (*Nadia keeps chickens*), **continue to be** if followed by an adjectival phrase (*Nadia kept calm*), and **continue to do** if followed by a gerund (*Nadia kept dancing*).

A second unused disambiguation cue is global context. Marker passing is used as a mechanism for local (intra-sentence) context cues, but the system has at present no representation of global context. It is my conjecture that it will not work simply to extend marker passing so that paths may be constructed between words of a sentence and the one before it. Rather, there should be a representation of context as an area in the knowledge base; this may include some nodes that were senses of words of previous sentences, instances created by the semantic representations of the sentences, nodes that participated in inferences made as sentences were read, and so forth. (Such representations of context are also motivated by the need to analyze reference and connections in discourse; see Hirst [1981a, 1981b].) Marker passing may then be extended to include this representation of context.

Because neither syntax nor global context are used yet, discourse focus cannot be used. Strictly speaking, Polaroid Words are wrong to even **try** to disambiguate *slug* in our example of section 3.3, *the slug operated the vending machine*. Rather, a clever system would have first recognized that the use of the definite article *the* implies that a disambiguated referent for the NP can be found in the focus, and no other disambiguation action need be taken (unless no referent is found). Of course, this wouldn't help if the NP were *a slug*.

The last unused cue is the requirement made by some verbs that certain of their cases must be present or that certain combinations of cases are prohibited. Adding this would allow preposition PWs to rule out possibilities in which none of them translate to a required case. In English, however, required cases are only a very weak cue, for English has few compulsory cases, and syntax serves to enforce most of them. A well-known example is the verb *break*, for which at least one of the cases AGENT, PATIENT, and INSTRUMENT must be present and be flagged by *SUBJ*. If we assume that the input is well-formed, then there will be a subject and it will be one of these three. An example of a compulsory case not enforced by syntax is the LOCATION case of the **place in position** sense of *put*:

(14) Ross put the luggage on the shelf.

(15) *Ross put the luggage.

This is a fact about the word *put* itself. An example of a prohibited combination of cases is the restriction for many (all?) verbs that an action in which the PATIENT is flagged by the surface subject may not have an INSTRUMENT expressed unless the verb is passive:

(16) The window broke.

(17) *The window broke with a rock.

(18) The window was broken with a rock.

4 What Polaroid Words can't do

It is possible, as I mentioned in section 3.2, that a PW could cross all its meanings off its list and suddenly find itself embarrassed. One possible reason for such an occurrence is that the word, or one nearby, is being used metaphorically, metonymically, or synecdochically, or, more generally, in a sense that the system does not know about. It is not possible in such cases to determine which word was actually responsible for the failure. Thus, if the **female swan** sense of *pen* is unknown, and a failure therefore occurs on (19):

(19) The pen flew ...

there is no way of telling that the missing meaning is in *pen* rather than *fly*. Ideally, the system should try to look for possible metaphors, as in (20):

(20) The pen flew across the page.

Research by Gentner [1981a, 1981b] suggests that if the system is looking for a possible metaphor, it should try the verb first, because verbs are inherently more "adjustable" than nouns; Gentner found that nouns tend to refer to fixed entities, while verb meanings bend more readily to fit the context. For example, people tend to paraphrase (21) [Gentner 1981a: 165]:

(21) The lizard worshipped.

as (22) rather than (23):

(22) The small grey reptile lay on a hot rock and stared unblinkingly at the sun.

(23) The nasty, despicable person participated in the church service.

Thus if a noun PW and a verb PW have irreconcilable differences, the noun should take precedence over the verb (regardless of which occurred first in the sentence—Gentner 1981a: 165).[16] If the system is still unhappy after doing its best to interpret the input metaphorically, it should ask the user for help and try again. None of this is included in present-day Polaroid Words.

Note that these problems occur only in cases where slot restrictions are tested. In the case of conflicting unambiguous words, one or both being used in a new, metaphoric, metonymic, or synecdochic sense, the conflict will not be noticed until the final Absity output is sent to Frail, since there is no reason to have checked for consistency. This will also be the case when strong marker-passing paths have caused a meaning to be chosen without checking slot restrictions, and in section 5.1 I show that this is a desirable state of affairs.

5 Psychological reality

Although cognitive modeling is a strategy for artificial intelligence system building rather than the main goal in this work, claims of psychological reality are interesting in their own right. In this section, therefore, I look at the degree of psychological reality in Polaroid Words with marker passing. I then discuss the importance of human data in the correct use of marker passing.

5.1 Polaroid Words, marker passing, and psycholinguistic models

Some degree of psychological reality was built into Polaroid Words from the start, in that their design was heavily influenced by the results of Swinney [1979; Onifer and Swinney 1981], who found that all meanings of an ambiguous word were accessed regardless of context. Similarly, PWs in all contexts start off with all

[16]There are, of course,exceptions to this general strategy. In particular, some sentences cannot be interpreted metaphorically, or any way other than literally; and sometimes the verb takes precedence over the noun. This sentence (due, I believe, to George Lakoff) exemplifies both cases: *My toothbrush sings five-part madrigals.* The word *madrigal* is quite unambiguous, and fits so well with the literal meaning of *sing*, that the incompatibility of selectional restrictions on `toothbrush` and the `agent` slot of `sing` is resolved in favor of the latter, and the sentence gives most people an image of a toothbrush that is somehow singing. [I am grateful to Eugene Charniak for pointing this out to me.]

their possible meanings available. This is in contrast to prior-choice models such as the Schank and Abelson [1977] script approach, in which one meaning has already been selected as the right one for the context, each context having its own lexicon to evoke. As we saw in section 1, the weight of evidence is now against these models. PWs also differ from ordered-search models in which meanings are tried one at a time, in a fixed order, until one fits.

Also, in accordance with Tanenhaus, Leiman, and Seidenberg's results [1979], all PW meanings are activated without regard for part of speech, since the marker passer has no access to syntactic information. Even though those meanings for what turns out to be the wrong syntactic category will be ignored, they will nevertheless have been used as the origin for marker passing, a fact which may affect later words (and probably adversely!—see section 5.2).

In addition, the use of marker passing as a mechanism for finding semantic associations was intended from the start to operate so as to model the effect of semantic priming. Unlike, for example, Hayes's CSAW disambiguation system [Hayes 1977a, 1977b], which only searches its network for associations when it needs them, the MP is continually active in the background, spreading markers around and finding and reporting associations even between unambiguous words. A node from which a marker has been passed is, of course, just like one that has been semantically primed by spreading activation, in that it has been set up to permit a rapid detection of concept association. (I will, however, qualify these remarks in the next section.) The separation of world knowledge from lexical knowledge, with marker passing occurring only in the former, suffices to prevent contextually inappropriate meanings of a word being pre-activated, in accordance with Lucas's results [1983].

However, the model is at variance with the results of Seidenberg, Tanenhaus, Leiman, and Bienkowski (*STLB*) [1982] published after the initial design of Polaroid Words (as described in Hirst and Charniak [1982]) was completed. STLB found that although multiple access is the norm, selective access seems to occur for semantically primed sentences such as (24):

(24) The farmer bought the straw.

All the same, a synthesis is possible if we make the following modification to PWs and MP. On starting, and before anything else happens, a new PW first checks whether any of its senses is already marked, that is, whether any immediate paths are available. If one is found, then that sense is chosen right away. Otherwise, things go on as before, with MP proceeding from all senses, looking for constructed paths. Since in the former case marker passing does not spread from unchosen meanings, they have not, in any real sense, been accessed. Thus, consistent with STLB's results, strong semantic priming would result in selective access but multiple access is still the norm—where strong semantic priming is, by definition, priming that results in an immediate marker passing path.

With or without this change, the model predicts speed differences in the disambiguation of semantically primed sentences such as (24), compared with non-priming biased sentences:

(25) The man walked on the <u>deck</u>.

The system will process (24) much faster than (25), because (24) should require only looking for MP chains, while (25) will look at the chains, find nothing, and then spend time dealing with slot restriction predicates in order to choose the right meaning of *deck*. The original model predicts two speed classes, the modified model predicts three. While these predictions seem plausible, there are as yet no relevant psychological data.

Polaroid Words also have the property that they can be led astray as people are by **negatively primed** sentences with **garden-path semantics** [Reder 1983]. Thus, PWs will make the same mistakes that most people do with sentences such as (26), in which the wrong meaning of an ambiguous word receives strong semantic priming:

(26) The astronomer married the <u>star</u>.

Reder found that comprehension 'errors'—that is, an inability to see the simple, literal meaning—increased markedly in such sentences. For example, in (26), people often take *star* to be **astronomical object** instead of **celebrity**, although this violates selectional restrictions on *marry*, and then become confused, or attempt to reinterpret the sentence metaphorically.[17] Similarly, MP will find the obvious chain between `astronomer` and `astronomical-object`; the PW for *star* will therefore choose the wrong meaning, and will not even

[17]There seem to be wide individual differences in the processing of such sentences. This is to be expected, as people will differ in their mental organization of concepts and hence in the exact boundary effects of spreading activation.

notice that it violates the slot predicate for `marry`, because it doesn't even consider such matters if it is happy with the MP path. The error will be discovered only after the sentence is fully interpreted and Frail attempts to evaluate the erroneous frame statement that was built. Similarly, intuition suggests that people who have trouble with the sentence only detect the error at a late stage, and then invoke some kind of conscious error recovery mechanism, such as interpreting the sentence as a metaphor.

However, many people do seem to be able to recover from disambiguation errors in which the garden-path pull is not as strong as that of (26). Daneman and Carpenter [1983: 566] gave subjects texts such as (27), in which subjects tended to initially choose the wrong meaning for the underlined word:

(27) The lights in the concert hall were dimmed. The audience watched intently as the famous violinist appeared on the stage. He stepped onto the podium and turned majestically to face the audience. He took a <u>bow</u> that was very gracefully propped on the music stand. The enthusiastic applause resounded throughout the hall.

They found that ability to recover from erroneous disambiguation correlated with scores on a reading span test,[18] which in turn correlated with verbal SAT scores.[19] Unfortunately, PWs presently recover from disambiguation errors like a reader with a very poor reading span—that is, they can't!

Lastly, PWs are in accord with psychological reality in that, when initially activated, they do not look at words to their right. They are thus in accord with the results of Stanovich and West [1983: 55], who found, contrary to previous suggestions [Underwood 1981], that there is in reading no effect from semantic characteristics of words in the right parafovea, that is, words to right of the fixation point.

In general, PWs (like the rest of the Absity system) are consistent with the trend of recent psychological results strongly suggesting that human language comprehension processes are modular, with limited interaction.

[18] In these tests, subjects had to read a set of sentences aloud, and then recall the last word of each. A subject's reading span is the size of the set they could handle without errors in recall.

[19] The Scholastic Aptitude Test (SAT) is a standardized test taken by applicants for college admission in the United States. It includes a test of verbal ability.

5.2 Psychological non-reality

In this section, I look at ways in which the system is at variance with psychological data, or for which psychological data are not yet available, but for which the predictions of PWs seem implausible [Hirst 1984b].

One important way in which MP differs from psychological reality is in the decay of spreading activation. The data of Tanenhaus, Leiman, and Seidenberg [1979] and Seidenberg, Tanenhaus, Leiman, and Bienkowski [1982] show that the facilitative effect of the activation of contextually incorrect meanings lasts less than 200 milliseconds (at least in those cases where rejection can be immediate). This suggests that activation of unchosen meanings decays very quickly. On the other hand, in the PW system all markers remain until the end of the sentence, at which time they are all reset. This may mean that paths are found between these unchosen senses and senses of later words, a clearly misleading occurrence. While the PWs (or path checker; see section 2.3) could check for such paths and eliminate them, it would be better if they didn't occur at all. At present, Frail does not support any form of selective removal of markers, so decay of activation from unchosen meanings could not be included in the present PW system.

The length of time that an ambiguous word may remain unresolved, with several alternatives active, probably also differs between PWs and people. In the present implementation, resolution will sometimes not take place until the very end of the sentence (see section 3.4).[20] Most psycholinguistic studies have looked only at cases where there is sufficient information to disambiguate a word as soon as it occurs, and the data on how long people will hold off resolution, hoping that more cues will come in if there is initially insufficient information, are equivocal.

The two studies on the question of which I am aware are that of Hudson and Tanenhaus [1984] and that of Granger, Holbrook, and Eiselt (*GHE*) [1984]. Hudson and Tanenhaus found that when there is no disambiguating information, both possible meanings of an ambiguous word remained active 500 msec after the word, but only one was active at the next clause boundary even though there had been no disambiguating information. The implication is that a best guess is made. GHE's findings were quite different. They looked

[20] An earlier decision may be forced by the Semantic Enquiry Desk if the information is needed to resolve a structural ambiguity; see Hirst [1987, sections 7.2.7 and 7.3.2].

at two-sentence texts in which the second sentence might require reinterpretation of an ambiguous word that had occurred, with misleading bias, in the first. For example:

(28) The CIA called in an inspector to check for bugs. Some of the secretaries had reported seeing roaches.

The first sentence is intended to make the reader decide that *bugs* means **hidden-microphones**, while the second requires that it actually be **insects**. After reading such sentences, subjects had to decide quickly whether a presented word was semantically related to the text. The error rate was extremely high, compared to control cases, when the word presented was related to the "incorrect" meaning of the ambiguous word of a sentence with misleading bias. GHE took this as evidence for **conditional retention**, *i.e.*, both meanings being retained for the ambiguous word, even across the sentence boundary.

The results of GHE are suspect for several reasons. First, their test for whether a word sense was active was not a proper on-line measure of activation. Second, their probe words are suspect. Thus for (28), their probes were *ant* (related to the text) and *spy* (unrelated to the text). Subjects who said that *spy* was related to the text were deemed to have made an "error", and this was taken as support for the hypothesis that the "incorrect" meaning of *bug* was still active. But clearly the word *spy* **is** related to the text, regardless of how the subject handles the word *bug*, simply because the word *CIA* was present in the text! Indeed, it is hard to imagine how to construct a probe that is sensitive to the activity of the "incorrect" meaning of the ambiguous word and yet not sensitive to the global context that was deliberately introduced to create the misleading bias.

What, then, can we reliably say about when a final decision is made on an ambiguous word? Intuition (which is not always a reliable source) suggests that while people will delay a decision on an ambiguous word for a little while, they will, nevertheless, usually make a final decision within a few words (or constituents?) of the ambiguity.[21] That a decision may be delayed is evidenced by the fact that people are not generally garden-pathed by sentences such as these:

(29) Nadia's favorite club is the five-iron.

[21] Nevertheless, there is evidence [Just and Carpenter 1980; Daneman and Carpenter 1983] for the psychological reality of a "sentence wrap-up" process in reading, in which remaining loose ends, such as unresolved references, are treated. It is possible that some residual disambiguation occurs as part of this process.

(30) Nadia's favorite <u>club</u> is The Carlton.

(31) Nadia's favorite <u>book</u> is *The House at Pooh Corner*. *(book =* **literary work***)*

(32) Nadia's favorite <u>book</u> is her signed first edition of *The House at Pooh Corner*. *(book =* **printed volume***)*

If people made an immediate "best guess" at the meaning of the ambiguous words *club* and *book*, then at least one of each of the above pairs should be inconsistent with the way a given individual is predisposed to guess, and therefore be a garden-path sentence for that individual. (Obviously, the strategy of making an immediate best guess would make language comprehension rather difficult at times.) It seems, therefore, that disambiguation of the examples above is incomplete until the final NP of the sentence is understood.

On the other hand, however, it is my intuition that a word such as *crook* in sentence (6) of section 3.4 is not resolved at the end of the sentence, as in PWs, but long before, possibly by the time the verb is processed:

(33) The crook operated ...

This choice seems to be based on probability: inanimate INSTRUMENT subjects are a lot less frequent than animate AGENTs, and, moreover, shepherd's staffs are rather unusual instruments for operating anything. The choice does not occur right after *crook*, on the basis of the relative infrequency of shepherd's staffs as a topic of modern urban conversation, as most people have no trouble with (34):

(34) The crook fell from the hook on which it was hanging.

This suggests that PWs should use a **cumulating evidence** approach and jettison unlikely alternatives quickly if there is no positive evidence for them. That is, one does not make an immediate best guess, but one does make a reasonable guess as soon as there is enough information to do so, even if one cannot be definite.[22] This has the advantage of helping to prevent combinatorial explosion. However, I have been loath to consider

[22]Kurtzman [1984] found that, in the case of structural ambiguity, the point of resolution varies widely, sometimes coming long before the disambiguating information, and sometimes not, in a manner very consistent with the idea of accumulating evidence. However, I know of no data addressing this issue for lexical ambiguity.

using this approach in Polaroid Words, in view of the dearth of data on the corresponding human behavior and the fuzziness of the whole notion. Any interim solution would have to fall back on “magic numbers”, and we have too many of those already (see next section). Nevertheless, PWs do use the relative frequency of the various meanings of an ambiguous word in some of their decisions (avoiding where possible tricks with magic numbers; see section 3.4), but since we know little of how people use word frequency in lexical ambiguity resolution, I have limited its use in PWs to tidying up loose ends at the end of a sentence. Another possibility that might be considered is to retain the present timing of decision making in PWs, but add a mechanism that watches out for looming combinatorial explosion and forces PWs to make an early guess if it senses danger.

Another prediction of Polaroid Words for which there is no psychological data is that case selection is performed by a subset of essentially the same mechanisms as lexical disambiguation. It is not clear how such a prediction could even be tested. The system also predicts the psychological reality of pseudo-prepositions, which almost certainly cannot be sustained.

Lastly, a few words should be said about marker passing. While we were happy to admit it as a discrete model of spreading activation, it should be pointed out that there are several competing theories of spreading activation. While these vary in their closeness to Frail marker passing, almost all of them differ in one important way. They assume that there is a certain limited amount of activation available to spread, and that spreading along a link ‘costs’ activation. Spreading stops when all activation is ‘spent’. When activation arrives at each node, it is divided among the outgoing arcs and an amount of it is passed along each. The division need not be equal; some links in the network may be **stronger** than others, and more activation spreads across strong links [Collins and Loftus 1975, Lorch 1982]. However, in Frail (and, hence, in MP) at present, all links are of equal strength, and all markers have the same status. (This is to be changed in future versions of Frail. It is also planned that the amount of activation spread from a node will be inversely proportional to the number of siblings it has [Hendler 1986].)

5.3 Marker passing, path strength, and magic numbers

One of the more vexed problems in using association cues for disambiguation is knowing when an association is strong enough to be considered conclusive evidence. We know from the existence of semantic garden-path sentences that associations alone should sometimes cause immediate jumping to a conclusion; we also know that this isn't true of all associations, for we do not get garden-pathed by sentences like (35):

(35) The lawyer stopped at the bar for a drink. *(*bar *is not taken in any of its legal senses)*

We therefore need some measure of the strength of an association, so that PWs will be able to jump to conclusions (rightly or wrongly) in the same situations that people do.[23] Although frequency of the use of the concepts should be a factor in determining the strength of an association [Anderson 1983],[24] I shall limit my remarks below to a discussion of the **semantic distance** between two concepts.

I mentioned in the previous section that most theories of spreading activation assume that different links have different strengths, though Frail does not attempt to model this. It is generally assumed that link strength is correlated with semantic distance—that a link between two concepts is strong exactly when they are very closely associated. Cases when this occurs may include one concept being a salient property of the other (**edibility**, **food**), or, possibly, a particularly good exemplar of the other (**robin**, **bird**);[25] a high frequency of use also strengthens a link [Collins and Loftus 1975, Anderson 1983] and hence the association between the concepts. On the other hand, de Groot [1983] has found that activation does not spread to associates of associates of a node—for example, **bull** and **cow** are linked and so are **cow** and **milk**, but activation from **bull** does not reach **milk**. Thus, PWs need a way to take an MP path and determine its strength, *i.e.*, the semantic distance between its endpoints, by looking at the links and nodes that it includes.

[23] We have already identified one such situation in section 5.1, namely, whenever an immediate MP path is found.

[24] Eugene Charniak (personal communication) has suggested that PWs should jump to a conclusion whenever marker passing selects a preferred meaning.

[25] It is often reported that people are faster at making categorization judgments for typical exemplars such as **robin–bird** than atypical ones such as **chicken–bird** [Rips, Shoben, and Smith 1973; Smith, Shoben, and Rips 1974; Collins and Loftus 1975]. This may be taken as evidence for the existence of typicality links, though Collins and Loftus [1975] show that it may be explained by the procedures by which positive and negative evidence for such decisions is gathered and evaluated.

The present, inadequate method of measuring path strength is a function of the length of the path, the nodes it passes through, and the links it uses. I use the following heuristics:

- The shorter the path, the stronger it is.
- The more arcs that leave a node, the weaker the connections through that node are (see the anti-promiscuity rule, section 2.3).

These methods, though I use them, are unsatisfactory because, like the marker passing constraints (section 2.3), they rely heavily on **magic numbers**. For example, the second suggests that any node will not be vague if it has only N arcs, but $N+1$ arcs will invariably tip the scale. This seems unlikely. And even if there were a neat threshold like that, how do we know that N is it?—it is merely a number that seems to work in the present implementation, but there is no principled reason for it. There is, of course, well-known evidence for the psychological reality of magic numbers in certain perceptual and short-term memory processes [Miller 1956], but it is hard to believe that this carries over to marker passing in long-term memory, where activation seems to be a continuous, not discrete, variable.

It is hoped that future versions of MP will be able to include such features as path strength and the weakening of activation as it gets further from the origin, so that we won't have to worry about post hoc measurements of path strength. This would be a first step in approximating the continuous nature of spreading activation.

6 Conclusion

I have presented a pair of cooperating mechanisms that both disambiguate word senses and determine case slots by finding connections between concepts in a network of frames and by negotiating with one another to find a set of mutually satisfactory meanings. PW processes work in parallel with a parser, Paragram, and a semantic interpreter, Absity, permitting them to deal with ambiguous words as if their semantic object were assigned immediately. (Hirst [1987, chapter 7] shows how PWs can also help in structural disambiguation.) The same PW control structure may be used for all syntactic categories. Polaroid Words minimize the need for separate, ill-motivated, purely linguistic knowledge; they have access to the system's world knowledge and use it wherever possible.

Polaroid Words are implemented as processes that interpret data structures containing purely lexical information that each word has in its dictionary entry. This is in contrast to approaches such as Small's [Adriaens and Small 1988], where the meaning of a word is represented as a large, barely-constrained procedure, different for every word, which parses and performs semantic interpretation as well as lexical disambiguation. Rather, the parser, Absity, and the marker passer do much of the work that Small requires his "word experts" to perform. We thereby capture generalizations in disambiguation, needing only one type of PW for each syntactic category and relying almost exclusively on general world knowledge.

Polaroid Words do not yet use syntactic disambiguation cues or global context, nor can they handle metaphor and metonymy. Table 1 shows some examples of sentences that the system can handle and some that it can't.

Acknowledgements

This work grew out of many discussions with Eugene Charniak. I have also received helpful comments on it from James Allen, Phil Hayes, Jim Hendler, Jennifer Holbrook, Susan Hudson, Margery Lucas, Susan McRoy, Amy Rakowsky, Anonymous Referee, and Nadia Talent. Jean-Pierre Corriveau and Diane Horton commented helpfully on earlier drafts of the paper. Parts of this paper are based on a chapter of Hirst [1987], and are used with the kind permission of Cambridge University Press. At Brown University, financial support for this work was provided in part by the U.S. Office of Naval Research under contract number N00014–79–C–0592 (Eugene Charniak, Principal Investigator). Preparation of this paper at the University of Toronto was supported by the Natural Sciences and Engineering Research Council of Canada.

References

Adriaens, Geert and Small, Steven L. (1988). "Word expert parsing revisited in a cognitive science perspective." in Small, Steven L.; Cottrell, Garrison W.; and Tanenhaus, Michael K. (editors). *Lexical ambi-*

Table 1: What Polaroid Words can and can't do

SENTENCES THAT CAN BE DISAMBIGUATED

SUBJ the slug operated *OBJ* the vending machine.

SUBJ the crook operated *OBJ* the pizza parlor.

SUBJ the crook wanted to kidnap *OBJ* Ross.

SUBJ Nadia's plane taxied to the terminal.

SUBJ the man walked on the deck.

SUBJ the deep philosopher threw *OBJ* the peach pit into the deep pit.

SENTENCES THAT CAN'T BE DISAMBIGUATED

The astronomer married the star.

Marker passing is misled.

The view from the window would be improved by a plant.

Requires inference.

I want to eliminate some moles.

No disambiguating information.

Ross was escorted from the bar to the dock.

Two parallel MP paths.

SUBJ the vending machine was operated by the slug.

No passives yet.

guity resolution: Perspectives from psycholinguistics, neuropsychology, and artificial intelligence. San Mateo, CA: Morgan Kaufmann Publishers, 1988. 13–43.

Anderson, John Robert (1983). "A spreading activation theory of memory." *Journal of verbal learning and verbal behavior*, **22**(3), June 1983, 261–295.

Becker, Joseph D. (1975). "The phrasal lexicon." in *Proceedings, [Interdisciplinary workshop on] Theoretical issues in natural language processing*, Cambridge, Massachusetts, June 1975. 70–73.

Charniak, Eugene (1983a). "A parser with something for everyone." in King, Margaret (editor). *Parsing natural language*. London: Academic Press, 1983. 117–149.

Charniak, Eugene (1983b). "Passing markers: A theory of contextual influence in language comprehension." *Cognitive science*, **7**(3), July–September 1983, 171–190.

Charniak, Eugene (1986). "A neat theory of marker passing." *Proceedings, Fifth National Conference on Artificial Intelligence (AAAI-86)*, Philadelphia, August 1986. 584–588.

Charniak, Eugene; Gavin, Michael Kevin; and Hendler, James Alexander (1983). "The Frail/NASL reference manual." Technical report CS–83–06, Department of Computer Science, Brown University, February 1983.

Collins, Allan M. and Loftus, Elizabeth F. (1975). "A spreading-activation theory of semantic processing." *Psychological Review*, **82**(6), November 1975, 407–428.

Daneman, Meredyth and Carpenter, Patricia A. (1983). "Individual differences in integrating information between and within sentences." *Journal of experimental psychology: Learning, memory and cognition*, **9**(4), October 1983, 561–581.

DeGroot, Annette M. B. (1983). "The range of automatic spreading activation in word priming." *Journal of verbal learning and verbal behavior*, **22**(4), August 1983, 417–436.

Fahlman, Scott Elliot (1979). *NETL: A system for representing and using real-world knowledge* (MIT Press series in artificial intelligence). Cambridge, Massachusetts: The MIT Press, 1979.

Gentner, Dedre (1981a). "Some interesting differences between nouns and verbs." *Cognition and brain theory*, **4**(2), Spring 1981, 161–178.

Gentner, Dedre (1981b). "Integrating verb meanings into context." *Discourse processes*, **4**(4), October–December 1981, 349–375.

Granger, Richard H., Jr; Holbrook, Jennifer K. and Eiselt, Kurt P. (1984). "Interaction effects between word-level and text-level inferences: On-line processing of ambiguous words in context." *Proceedings, Sixth Annual Conference of the Cognitive Science Society*, Boulder, June 1984. 172–178.

Hayes, Philip J. (1977a). *Some association-based techniques for lexical disambiguation by machine.* Doctoral dissertation, Département de Mathématiques, École polytechnique fédérale de Lausanne. Published as technical report 25, Department of Computer Science, University of Rochester, June 1977.

Hayes, Philip J. (1977b). "On semantic nets, frames and associations." *Proceedings, 5th International Joint Conference on Artificial Intelligence*, Cambridge, Massachusetts, August 1977. 99–107.

Hendler, James Alexander (1986). *Integrating marker-passing and problem-solving: A spreading-activation approach to improved choice in planning.* Doctoral dissertation [available as technical report CS–86–01], Department of Computer Science, Brown University, January 1986.

Hendler, James Alexander (1987). "Issues in the design of marker-passing systems." in Boudeaux, J. C.; Hamill, Bruce W.; and Jernigan, Robert (editors). *The role of language in problem solving 2.* Amsterdam: North-Holland, 1987, 227–245.

Hirst, Graeme (1981a). *Anaphora in natural language understanding: A survey* (Lecture notes in computer science 119). New York: Springer-Verlag, 1981.

Hirst, Graeme (1981b). "Discourse-oriented anaphora resolution in natural language understanding: A review." *American journal of computational linguistics*, **7**(2), April–June 1981, 85–98.

Hirst, Graeme (1983). "A foundation for semantic interpretation." *Proceedings, 21st Annual Meeting of the Association for Computational Linguistics*, Cambridge, Massachusetts, June 1983. 64–73.

Hirst, Graeme (1984a). "A semantic process for syntactic disambiguation." *Proceedings, Fourth National Conference on Artificial Intelligence (AAAI-84)*, Austin, August 1984. 148–152.

Hirst, Graeme (1984b). "Jumping to conclusions: Psychological reality and unreality in a word disambiguation program." *Proceedings, Sixth Annual Conference of the Cognitive Science Society*, Boulder, June 1984. 179–182.

Hirst, Graeme (1987). *Semantic interpretation against ambiguity* (Studies in natural language processing). Cambridge, England: Cambridge University Press, 1987.

Hirst, Graeme and Charniak, Eugene (1982). “Word sense and case slot disambiguation.” *Proceedings, National Conference on Artificial Intelligence*, Pittsburgh, August 1982. 95–98.

Hudson, Susan B. and Tanenhaus, Michael K. (1984). “Ambiguity resolution in the absence of contextual bias.” *Proceedings, Sixth Annual Conference of the Cognitive Science Society*, Boulder, June 1984. 188–192.

Just, Marcel Adam and Carpenter, Patricia A. (1980). “Inference processes during reading: From eye fixations to comprehension.” *Psychological review*, **87**(4), July 1980, 329–354.

Kurtzman, Howard Steven (1984). *Studies in syntactic ambiguity resolution*. Doctoral dissertation, Department of Psychology, Massachusetts Institute of Technology, 13 September 1984. Indiana University Linguistics Club.

Lorch, Robert F., Jr (1982). “Priming and search processes in semantic memory: A test of three models of spreading activation.” *Journal of verbal learning and verbal behavior*, **21**(4), August 1982, 468–492.

Lucas, Margery M. (1983). “Lexical access during sentence comprehension: Frequency and context effects.” *Proceedings, Fifth annual conference of the Cognitive Science Society*, Rochester, New York, May 1983. [unpaginated]

Marcus, Mitchell P. (1980). *A theory of syntactic recognition for natural language.* Cambridge, MA: The MIT Press, 1980.

Miller, George Armitage (1956). “The magical number seven, plus or minus two: Some limits on our capacity for processing information.” *Psychological review*, **63**(2), March 1956, 81–97.

Onifer, William and Swinney, David A. (1981). “Accessing lexical ambiguities during sentence comprehension: Effects of frequency of meaning and contextual bias.” *Memory and cognition*, **9**(3), May 1981, 225-236.

Quillian, M. Ross (1962). "A revised design for an understanding machine." *Mechanical translation*, **7**(1), July 1962, 17–29.

Quillian, M. Ross (1968). "Semantic memory." in Minsky, Marvin Lee (editor). *Semantic Information Processing.* Cambridge, Massachusetts: The MIT Press, 1968. 227–270.

Quillian, M. Ross (1969). "The teachable language comprehender: A simulation program and theory of language." *Communications of the ACM*, **12**(8), August 1969, 459–476.

Reder, Lynne M. (1983). "What kind of pitcher can a catcher fill? Effects of priming in sentence comprehension." *Journal of verbal learning and verbal behavior*, **22**(2), April 1983, 189–202.

Rips, Lance J.; Shoben, Edward J.; and Smith, Edward E. (1973). "Semantic distance and the verification of semantic relations." *Journal of verbal learning and verbal behavior.* **12**(1), February 1973, 1–20.

Schank, Roger Carl and Abelson, Robert Paul (1977). *Scripts, plans, goals and understanding: An enquiry into human knowledge structures.* Hillsdale, New Jersey: Lawrence Erlbaum Associates, 1977.

Seidenberg, Mark S.; Tanenhaus, Michael K.; Leiman, James Mehner; and Bienkowski, Marie A. (1982). "Automatic access of the meanings of ambiguous words in context: Some limitations of knowledge-based processing." *Cognitive psychology*, **14**(4), October 1982, 489–537.

Smith, Edward E.; Shoben, Edward J.; and Rips, Lance J. (1974). "Structure and process in semantic memory: A featural model for semantic decisions." *Psychological review*, **81**(3), May 1974, 214–241.

Stanovich, Keith E. and West, Richard F. (1983). "The generalizability of context effects on word recognition: A reconsideration of the roles of parafoveal priming and sentence context." *Memory and cognition*, **11**(1), 1983, 49–58.

Swinney, David A. (1979). "Lexical access during sentence comprehension: (Re)Consideration of context effects." *Journal of verbal learning and verbal behavior*, **18**(6), December 1979, 645–659.

Tanenhaus, Michael K.; Leiman, James Mehner; and Seidenberg, Mark S. (1979). "Evidence for multiple stages in the processing of ambiguous words in syntactic contexts." *Journal of verbal learning and verbal behavior*, **18**(4), August 1979, 427–440.

Underwood, Geoffrey (1981). "Lexical recognition of embedded unattended words: Some implications for reading processes." *Acta psychologica*, **47**, 1981, 267–283.

Wilensky, Robert and Arens, Yigal (1980a). "PHRAN: A knowledge-based approach to natural language analysis." Memo UCB/ERL M80/34, Electronics Research Laboratory, University of California, Berkeley, 12 August 1980.

Wilensky, Robert and Arens, Yigal (1980b). "PHRAN: A knowledge-based natural language understander." *Proceedings, 18th Annual Meeting of the Association for Computational Linguistics*, Philadelphia, June 1980. 117–121.

Wong, Douglas (1981a). "Language comprehension in a problem solver." *Proceedings, 7th International Joint Conference on Artificial Intelligence*, Vancouver, August 1981. 7–12.

Wong, Douglas (1981b). *On the unification of language comprehension with problem solving.* Doctoral dissertation [available as technical report CS-78], Department of Computer Science, Brown University, 1981.