9 Speculations, Partially Baked Ideas, and Exercises for the Reader

> When there is light at the end of the tunnel, order more tunnel.
> —Anonymous[1]

> But in our enthusiasm, we could not resist a radical overhaul of the system, in which all of its major weaknesses have been exposed, analyzed, and replaced with new weaknesses.
> —Bruce Leverett[2]

The research I have described in this book necessarily leaves unanswered many questions, big and small. Many sections, especially 3.8, 5.3.6, 5.4, and 7.4 have discussed things left undone by Absity, Polaroid Words, and the Semantic Enquiry Desk. In this chapter, I list a number of other open questions, partially baked ideas, and wild speculations, sometimes with my thoughts on how they might be answered or developed. Some could be dissertation topics; others may be good subjects for a term paper or course project. Several are psycholinguistic experiments. At the start of each question (if appropriate) I give in brackets the section or sections of this book in which the matter is discussed.

9.1 The representation of knowledge

Exercise 1.1 [1.1.2, 1.3.1, 5.2, 5.6.3] Could a non-discrete representation of knowledge be developed for AI? Such a representation would be able to handle close similarities and differences, such as the **head of a pin** compared with the **head of a hammer**. Consider the possibility of pseudo-continuous representations that are to discrete representations as floating-point numbers are to integers. Candidates

[1] Found in the 'fortune' database distributed with 4.1c BSD UNIX. *UNIX* is a trademark of AT&T Bell Laboratories.

[2] LEVERETT, Bruce W. *Register Allocation in Optimizing Compilers*. Doctoral dissertation [available as technical report CMU–CS–81–103], Department of Computer Science, Carnegie–Mellon University, February 1981. 134.

to consider include some kind of network of neuron-like nodes (*cf.* Feldman and Ballard 1982; Feldman 1985), a value-passing machine such as Fahlman, Hinton and Sejnowski's (1983) Thistle system, and a simulated-annealing or Boltzmann network (Kirkpatrick, Gelatt and Vecchi 1983; Fahlman, Hinton and Sejnowski 1983; Smolensky 1983) *(see also exercise 4.8)*. A concept would be represented by a very large number of nodes and connections instead of just a few as in Frail, and no one node would be essential to the representation of any concept. Two concepts would be similar to the extent that they incorporated the same nodes.

How could Absity and its friends operate with such a representation? For example, in lexical disambiguation, how could a relationship between two concepts, such as **astronomer** and **star**, be found? Can marker passing operate between groups of nodes? What about relationships implied by two concepts incorporating the same node?—such relationships are not necessarily relevant to lexical disambiguation.

Exercise 1.2 [5.2.3] Despite all efforts to the contrary, some of the paths found by a marker passer will be uninteresting or misleading. Give an example of such a path that could not have been prevented without also losing desirable paths. Charniak (1982, 1985) has proposed that the output of a marker passer should be filtered though a path checker that would attempt to remove undesirable paths. Characterize as formally as possible what characteristics make a path undesirable. Can undesirable paths be categorized into easy-to-recognize classes?

Exercise 1.3 [7.3.2] The SED's gap-finding methods ideally require that Frail be able to take a frame statement and search for "something similar". This requires a better and more formal idea than we presently have of what similarity of instances is. Can you develop such a formalization? To begin, consider that if the frame statement to be matched does not describe an extant instance (which will usually be the case, or there wouldn't be a need to search for a similarity), it can be given a "phantom instantiation" and thereupon treated as extant. Then marker passing from the phantom could be a simple way to find candidates for similarity, though each candidate would have to be tested, as even careful marker passing (presumably with constraints different from those used to find simple relationships by Polaroid Words) would probably find a large number of spurious candidates. Would marker passing be adequate for finding candidates? Or would it be common for instances to be distant from others to which they are similar? Techniques such as Winston's (1978) transfer frames may also be applicable to this problem.

Exercise 1.4 [4.3, 4.3.1] **Literary criticism made simple:** Many styles of poetry operate simply by forcing the user to make a mental connection that he or she did not previously have, viewing the world in a slightly different way. When the poet succeeds in making a felicitous long-distance connection in the user's head, the latter experiences a 'delight' response. This is best seen in minimalist poems such as those of Richard Brautigan. Consider, for example, these poems:

(9-1) **The pill versus the Springhill mine disaster**

When you take your pill
it's like a mine disaster.
I think of all the people
lost inside of you.[3]

(9-2) **The net wt. of winter is 6.75 ozs.**

The net wt. of winter is 6.75 ozs.
and winter has a regular flavor
with Fluoristan to stop tooth decay.

A month ago I bought a huge tube
of Crest toothpaste and when I put it
in the bathroom, I looked at it
and said, "Winter."[4]

(9-3) **Romeo and Juliet**

If you will die for me,
I will die for you

and our graves will
be like two lovers washing
their clothes together
in a laundromat.

If you will bring the soap,
I will bring the bleach.[5]

(9-4) **Kafka's hat**

With the rain falling
surgically against the roof,
I ate a dish of ice cream
that looked like Kafka's hat.

It was a dish of ice cream
tasting like an operating table
with the patient staring
up at the ceiling.[6]

---

[3]BRAUTIGAN, Richard. "The pill versus the Springhill mine disaster." *The pill versus the Springhill mine disaster*, New York: Dell, 1968. 100. Copyright © 1968 by Richard Brautigan. Reprinted by permission of Delacorte Press / Seymour Lawrence and Helen Brann Agency, Inc.

[4]BRAUTIGAN, Richard. "The net wt. of winter is 6.75 ozs." *Rommel drives on deep into Egypt*, New York: Delta, 1970. 12. Copyright © 1970 by Richard Brautigan. Reprinted by permission of Delacorte Press / Seymour Lawrence and Helen Brann Agency, Inc. *Crest* is a trademark of Proctor and Gamble.

[5]BRAUTIGAN, Richard. "Romeo and Juliet." *Rommel drives on deep into Egypt*, New York: Delta, 1970. 7. Copyright © 1970 by Richard Brautigan. Reprinted by permission of Delacorte Press / Seymour Lawrence and Helen Brann Agency, Inc.

[6]BRAUTIGAN, Richard. "Kafka's hat." *The pill versus the Springhill mine disaster*, New York: Dell, 1968. 89. Copyright © 1968 by Richard Brautigan. Reprinted by permission of Delacorte Press / Seymour Lawrence and Helen Brann Agency, Inc.

The reader probably did not previously have a mental connection between toothpaste and winter; the poem's purpose is to make one. Whether or not such poems 'work' for a particular reader must depend on whether that person's knowledge base has a structure that permits the association to be made, and it is probably the case that for any one poem, a hit rate of at best about 25% can be expected—the other 75% of readers won't get it. (I hope that at least one of the above examples worked for you.)

Exercise: Write a Brautigan appreciater—a program very willing to make connections. Extend it to handle Eliot or Yeats or Leonard Cohen. For extra credit: Write a program that generates Brautigan-like poems. (Do not attempt to extend it to handle Eliot or Yeats or Leonard Cohen.)

Wax and gold—Literary criticism made difficult: Donald Levine (1965, 1985: 21–28) describes the importance of deliberate linguistic ambiguity in a large number of non-Western cultures. He writes that, for example,[7]

> in one African variant of Islamic culture, that of the Somali nation, a love for ambiguity appears particularly notable in the political sphere. David Laitin reports that the Somali boast that "the Somali language is sinuous", because it permits words to take on novel shapes that accommodate a richness of metaphors and poetic allusions. Political arguments and diplomatic messages take the form of alliterative poems, mastery of which is a key to prestige and power. These poems typically begin with long, vague, circumlocutory preludes, introducing the theme at hand, which is then couched in allegory. Of these poems, Laitin writes:
>
> > A poetic message can be deliberately misinterpreted by the receiver, without his appearing to be stupid. Therefore, the person for whom the message was intended is never put in a position where he has to answer yes or no, or where he has to make a quick decision. He is able to go into further allegory, circling around the issue in other ways, to prevent direct confrontation. (Laitin 1977: 39)

Levine is particularly struck by the Amhara culture of Ethiopia.

> In what is perhaps the most characteristic expression of the Amhara genius, a genre of oral literature known as "wax and gold", the studied use of ambiguity plays a central part. Wax and gold (*sam-ennā warq*) is the formula with which the Amhara symbolize their favorite form of verse. The form consists of two semantic layers. The apparent, superficial meaning of the words is called "wax" (*sam*); their hidden, deeper significance is the "gold" (*warq*).
>
> The following Amharic couplet exemplifies the *sam-ennā warq* figure:
>
> Since Adam your lip did eat of that Tree
> The Savior my heart has been hung up for thee.

[7] The following excerpts from Levine 1985: 23–27 and Laitin 1977: 39 are copyright © 1977, 1985 by The University of Chicago, by whose kind permission they are reprinted here.

In this secular couplet the "wax" of Adam's sin and Christ's crucifixion in his behalf has been used as a mold in which to pour a love message. A literal translation of the "wax" of the couplet is:

> Because Adam ate of the apple from the Tree of Knowledge
> The Savior of the World has been crucified for thee.

To appreciate the "gold" of the couplet, one must know that the verb meaning "was crucified", *tasaqqala*, may also mean "is infatuated with". A literal translation of the "gold" content would be:

> Because of your (tempting) lips
> My heart is infatuated with thee.

In other figures, the duplicity of the message is rendered less explicit. In figures known as *hĭber* and *merĭmer*, the "wax" and "gold" are combined in the same word or phrase instead of being put side by side. These figures thus correspond to the English pun. For example:

> Your father and your mother have vowed to keep from meat
> But you, their very daughter, innards do you eat.

"To eat someone's entrails" is an Amharic idiom which means "to capture his heart". The hidden meaning of the couplet is thus: "You made me love you".

Exercise: Write a wax-and-gold appreciater—a program that can find both the surface and hidden meanings of a text. The program should also appreciate puns in English.

Exercise 1.5 [8.2.3] Must a knowledge representation for medical diagnosis necessarily be different from the compositional frame-like representations we require as a target for semantic interpretation? If so, how can people learn or amend their rule bases from reading? Is a synthesis of the two kinds of representation possible? (Recall that Frail itself was developed as a synthetic representation for both language understanding and problem solving.)

Exercise 1.6 [5.5] I have argued that selectional restrictions are better accounted for as slot restriction predicates in the knowledge base than as symbols on words in the lexicon.[8] This may be a slight oversimplification, however; that is, there may be words that are essentially synonymous, pointing to the same frame in the lexicon, and yet have different selectional restrictions. For example, *eat* and the intransitive sense of *feed* are probably both best represented by the `eat` frame, yet *feed*, unlike *eat*, requires its AGENT to be an animal:

(9-5) Cows feed on hay.

(9-6) #Ross feeds on pancakes.

[8] Lytinen (1984: 33) also argues for this.

Sentence (9-6) seems slightly metaphorical, implying that Ross is voracious or animal-like in his approach to pancakes.[9]

Our approach can be saved by making distinct frames for *eat* and *feed*; but is this reasonable? The fact that some other languages make the same distinction (*e.g.*, *essen* and *fressen* in German) suggests that it may be—indeed, it may be necessary in an interlingua for machine translation. What other possible counterexamples are there? Could a system with separate lexical selectional restrictions and slot restriction predicates be feasible? How should one handle the restrictions on words like *board*, where one can only use *board* for a vehicle that one can stand in or on—a boat, airplane, hovercraft, or bus, for example, but not a canoe or a car?

Exercise 1.7 [1.3.1, 3] We have shown Frail to be a reasonable target for semantic interpretation. However, for many applications, such as machine translation, this is insufficient; it must also be a suitable starting point for language generation. What are the necessary properties? Does Frail have them?

9.2 Semantic formalisms

Exercise 2.1 [2.2.1, 3.8] Although we have been very critical of decompositional semantics, frame-based semantics is vulnerable to many of the same criticisms. The frame to define a chair, for example, is little more than a structured version of the decompositional representation in (2-8); the symbols have now become slots, frames, and ISA relations. Still remaining are such problems as how an inherently vague concept such as **chair** can be adequately defined and what a suitable set of primitives is. The problem of choosing primitives is especially important in systems like Frail that, unlike, say, KRL *(see section 1.2)*, maintain a strict ISA hierarchy without overlap of subclasses. Investigate possible solutions to these problems. How can inherent vagueness be handled in a frame system? Methods to discuss include defining a concept as a procedure that tests whether an object meets the necessary criteria or as a prototype with tolerances for variation from it (*cf.* Winograd 1976: 13–14), and using the mechanisms of fuzzy logic (Zadeh 1983). Remember that we wish to retain the full inference power of frames and our ability to use them as compositional semantic objects. Compare network frame systems and strictly hierarchical frame systems for expressive power, deductive abilities, and implementation efficiency.

Exercise 2.2 [2.1] A new approach to meaning, SITUATION SEMANTICS, has recently been made popular by Barwise and Perry (1983; Barwise 1981; Israel 1983;

[9] Transitive uses of *feed* do not follow this pattern; the feedee may be human or animal:

(i) Nadia fed the geese corn.

(ii) Ross fed Nadia pancakes.

Even so, sentence (ii) seems to suggest that the situation lacked a certain gentility, for otherwise *served* would have been a more natural verb to use.

Cooper 1985). Situation semantics is a REALIST theory, in that it takes meaning to be something actually in the world, and linguistic meaning to be simply a special case of a more general phenomenon. Absity is based on the Montague (1973) way of thinking about semantics; could situation semantics similarly form the basis for a compositional computational semantic formalism (*cf.* Winograd 1984)? What would the resulting semantic interpreter be like? What would its advantages and disadvantages be compared to those of Absity? (For some work on these questions, see Fenstad, Halvorsen, Langholm, and van Bentham 1985, and Lespérance 1986.)

Exercise 2.3 [2.2.2, 3.3] Chierchia (1982, 1983) has pointed out certain problems with maintaining consistency in Montague's (1973) typing system. For example, the NP *Ross* and the infinitive verb *to run* have different types—*individual* and *property of individual*, respectively—which in turn obliges *is nice* to be of different types in (9-7) and (9-8), a clearly undesirable and counterintuitive effect:

(9-7) Ross is nice.

(9-8) To run is nice.

Chierchia's solution is a less prolific, three-level system of types in which properties such as *to run* may be used interchangeably with individuals such as *Ross*. The typing of Absity, also much simpler than Montague's, is able to handle the particular example above, because an NP is a frame statement, and a verb is a frame (which may be turned into a frame statement by combining it with the determiner *NULLDET* to make an NP—though to deal with (9-8) properly, its intensionality must also be handled). But Absity may have consistency problems nevertheless. Can you find such a case? Could consistency be dependent upon the particular knowledge representation? If so, could Chierchia's solution be adapted for use in Absity? What are the ramifications for Frail?

Exercise 2.4 [2.2.2, 3.3; continues previous exercise] Just as Montague semantics makes counterintuitive type distinctions *(see previous exercise)*, it also fails to make distinctions that we intuitively feel it ought (Chierchia 1982, 1983). For example, in (9-9) and (9-10), *slowly* and *try to* both turn out to be functions from properties to properties:

(9-9) Ross eats the ice cream slowly.

(9-10) Ross tries to eat the ice cream.

Since Absity maps different syntactic types to the same semantic type (*e.g.*, adverbs and auxiliaries are both slot–filler pairs), it is vulnerable to the same criticism. Show that Absity's type mappings are intuitively well motivated. A fortiori, show that their motivation is not unlike that for Chierchia's approach to types *(see previous exercise)*.

Exercise 2.5 [3.2] Is there a principled way in which it may be decided which rules of a grammar should have semantic rules associated with them?

Exercise 2.6 [3.8, 7.2.1, 7.2.6] Frail has no way of representing counterfactual states, as in (9-11):

(9-11) Nadia thinks of her house with a new coat of paint.

The underlined phrase describes an instance that is the same as an extant instance except for a slot value or two but has only conceptual existence, though it doesn't seem to be merely an intension. Implement a suitable representation for such things, and show how your representation can be used to determine attachment of non-restrictive prepositional phrases to uniquely defined noun phrases, as in (9-11).

9.3 Semantic interpretation and discourse pragmatics

The Absity system derives the literal meaning of its input, and is almost completely devoid of knowledge of discourse pragmatics. The following exercises address these deficiencies.

Exercise 3.1 Polaroid Words provide a suitable place to handle finding the referents of pronouns. The PWs for words such as *she*, *it*, and *they* could externally appear to be like an ordinary noun Polaroid Word. Internally, they could use whatever methods were appropriate (Hirst 1981a, 1981b) to determine the set of candidate referents, and then be disambiguated like any noun PW (*cf.* Pazzani and Engelman 1983; *see section 8.3.1*). Non-pronominal definite references require a different approach: appropriate extension of the abilities of frame determiners in Frail. This gives the appearance of having two separate reference resolution mechanisms, but this need not be the case; internally, the PW for a pronoun may use frame determiners. For example, the PW for *she* could behave just like the Frail call `(the ?x (female ?x))`. Implement these additions to Polaroid Words and Frail. You will need to be able to determine which of the concepts previously mentioned in the discourse are candidate referents and which in particular have been highlighted by the speaker or writer as the discourse topic. See if the various methods described in Hirst 1981a, 1981b can be adapted for use with Frail.

Exercise 3.2 [5.3.3] Neither Absity nor the Polaroid Word system has any method for dealing with the various categories of "canned phrase" in English, except that in section 5.3.3 we conceptually incorporated by fiat the "phrasal recognizer" of Wilensky and Arens (1980a, 1980b). Such an incorporation will not necessarily be straightforward. The more rigid expressions, especially ungrammatical ones such as *by and large*, could well be handled with a little lookahead at lexicon lookup time and replaced by a single token in the input stream; this would fit well with the present system. On the other hand, some canned phrases, such as *kick the bucket* (= **die**), may also be taken literally or be subject to syntactic and morphological processes; these cannot be handled at the initial lexicon-lookup level, but can only be recognized after some parsing and interpretation. This could lead to control and interaction problems with both Absity and Polaroid Words. For example, the system will faithfully interpret (9-12) literally:

(9-12) Ross kicked the ...

and, upon finding the next word to be *bucket*, will first have to ask whether Ross is really kicking a real bucket or not; if he isn't, it will then have to get rid of its partially built literal interpretation. Could Absity handle this? Consider also Weischedel's (1983) method of using Horn clauses to map from one semantic representation of a sentence to another (in this case from its literal to its intended meaning).

A slightly different approach, however, is suggested by psycholinguistic evidence that people do not process idioms or canned phrases in such a two-stage process, trying first the literal meaning and then the idiomatic one. Rather, Ortony, Schallert, Reynolds, and Antos (1978) found that idioms were processed as fast as control phrases and that processing was slowest when the context indicated the literal meaning of the canned phrase. It seems likely that canned phrases are stored as a single unit in the mental lexicon and accessed very early; thus, *cross that bridge* would be enough to cause access to the phrase *cross that bridge when one comes to it*. Similarly, McDonnell (1982) found that an appropriate context can prime the meaning of idioms, just as ordinary words can be primed. However, Rakowsky (1985) has found that idioms themselves do not prime their idiomatic senses, and, further, seem to inhibit ("negatively prime") their literal senses. Rakowsky suggests a model in which processing of an idiomatic sense starts as soon as the sense can be accessed. Processing of the literal sense, however, remains incomplete, but a trace of it remains so that it may be completed if it turns out to be needed.

It should be possible to do something similar in Absity. A process would look out for canned phrases, and when it found one it would take over from Absity for the duration of the idiom (but saving Absity's state, just in case) and insert a structure with the appropriate meaning in the interpretation. The structure would be Polaroid Word–like to the extent that it would detect whether or not it fitted properly; if it found, sooner or later, that it didn't, it would restore Absity. This of course makes semantic interpretation no longer deterministic, which might be regarded as too drastic.

Exercise 3.3 [5.4] How may metaphor be detected and resolved by Polaroid Words?

Exercise 3.4 [continues previous exercise] Not all constructs that violate slot restriction predicates are metaphoric. Consider:

(9-13) Ross believes that rocks can fly.

(9-14) Nadia doubts that rocks can fly.

Neither meaning of *rocks*, **stones** and **rocking movements**, fits with *fly*, but of course the sentences are perfectly acceptable and **stones** is the sense chosen by informants. This could be simply because it is the more frequent sense, or it could be that, as a **physical object**, it is still the better choice for **fly**ing. Experiment: Determine which.

The problem for Polaroid Words here is deciding when to relax slot restrictions and allow a literal interpretation. In the examples above, we might say that the context, a belief report, gives "permission" for this, just as children's story or fantasy may permit talking animals and the like. Matters are not this simple, however, or else we couldn't use metaphors in belief reports:

(9-15) Nadia believes that computers are the rocks in the bowling green of life.

More importantly, there is simply a lot of nonsense in the real world:

(9-16) **Go on, have your fun. It's always the children that suffer later:** Los Angeles secretary Jannene Swift married a 50-pound pet rock in a formal ceremony in Lafayette Park.[10]

The context is strong enough to force the reading that Ms Swift went through the motions of marriage, even if the preconditions were not fulfilled and there is no marriage, legally or semantically. Again, it may help that *rock* can be a **physical object**;[11] it is hard to assign a literal interpretation, no matter how anomalous, to (9-17):

(9-17) Ross was married to his career.

Exercise: Investigate the conditions under which anomalous sentences are and aren't interpreted figuratively.

Exercise 3.5 Subtleties of meaning are conveyed in a speaker's or writer's exact choice of words (Lanham 1974); to translate (9-18) and (9-19) into effectively the same representation in Frail would be to miss this point:

(9-18) We cannot dedicate—we cannot consecrate—we cannot hallow—this ground. The brave men, living and dead, who struggled here, have consecrated it far above our poor power to add or detract.[12]

(9-19) Some men who were brave, some of whom are now dead, and who were faced with much adversity at this particular place have already gone and made said place holy, thereby effectively preventing us from doing it now.

How could the difference between (9-18) and (9-19) be detected by Absity and represented in Frail?

Exercise 3.6 [6.2.1, 7.4.3] Discourse comments, such as *frankly* in (9-20):

(9-20) Frankly, it gives me a headache.
(i.e., I am frank in saying this: it gives me a headache.)

[10] "Esquire's dubious achievement awards for 1976." *Esquire*, **87**(1), January 1977, 49–55.

[11] It also helps that *pet rock* has become an unambiguous canned phrase. When it was novel, however, interpreting it involved exactly the same problems of trading off metaphor for *pet* and literal interpretation for *rock*.

[12] LINCOLN, Abraham. 19 November 1863.

are not part of the literal interpretation of a sentence but operate at a different level. How may they be detected and correctly handled by Absity? Consider also the related problem of EPITHET INSERTION, in which a word is inserted whose literal interpretation should be ignored, its presence instead serving as an intensifier:

(9-21) I fucking hate this fucking assignment![13]

Exercise 3.7 [3.8, 8.2.3, exercise 1.5] An interlingua for machine translation must represent much more than literal meaning. Other necessities include non-literal meaning (*i.e.*, intent of speaker), the topic or emphasis, and the tenor or tone. A sentence representation thus becomes an n-tuple, of which the Frail semantic representation is just one component. We would like the processes that produce the other components to be as closely related to Absity as possible. Discuss the architecture of such a system.

Exercise 3.8 [3.8] Make a full inventory of words such as *again*, *even*, and *let alone* which are problematic for compositional semantics. What patterns do you see? Can Absity be improved to deal with these words, possibly by treating them as functions instead of passive objects? If you believe that a separate pragmatic process is necessary to handle them, design such a process. How can the import of such words be carried through a language-independent interlingua in a machine translation system? Make a cross-linguistic study of these words, and look for universals.

9.4 Lexical ambiguity

Exercise 4.1 [1.1.2] The Boots-And-All Theory of Language Comprehension (Hirst 1981a[1]:49, footnote 28) says that language tends to evolve so that every available cognitive faculty is used in its comprehension. Therefore, if there were a real cognitive distinction between polysemy and homonymy, as Panman's work (1982) suggests, we would expect that at least some languages, if not all of them, would exploit the distinction in some way, and that this would be reflected in any formal description of such a language. Find such a language, or, better still, show that English is such a language. Alternatively, show that a language could make use of the distinction without it being "noticeable" in a formal description.

Exercise 4.2 Some words are ambiguous at the morphological level. For example, an *undoable knot* could be one that can be undone or one that cannot be done—that is, *undo-able* or *un-doable*.[14] It is probably not adequate simply to mark all such words in the lexicon, as affixes such as *un-* and *-able* are highly productive. Examine the size of the class of words for which this is a problem,

[13] Jim Hendler, personal communication, 22 September 1983.

[14] I am grateful to Martin Kay for this example.

and see if simple lexical ambiguity really has to be ruled out. Could novel words of this class be reliably identified by Polaroid Words in order that they be given special treatment? What treatment?

Exercise 4.3 [1.1.1] In English, ambiguous verbs tend to be polysemous, while ambiguous nouns tend to be homonymous; adjectives show less ambiguity than nouns and verbs. Is this true in other languages? Run an experiment to test the hypothesis that polysemy is "cognitively easy" for verbs and homonymy "cognitively hard", while the reverse is true for nouns.

Exercise 4.4 If a polysemous PW is unable to choose between two or more of its meanings, one strategy it could use is to take the "central concept" common to these senses (*cf.* Marslen-Wilson and Tyler 1980). In a frame system such as Frail, this central concept might be the nearest frame dominating both candidate senses in the ISA hierarchy. Test this suggestion by implementing it. Note the trade-offs in deciding whether to dynamically search for this frame when it is needed, or to store it in the lexical entry for the PW; the latter case, though faster, would be difficult if the frame sought depended on exactly which senses remained, and would be awkward if the frame system were frequently changing. In either event, the PW would have to be marked to indicate whether use of the strategy is permitted, since applying it to a homonym could be disastrous.

Exercise 4.5 [4.2.4, 4.3, 4.3.3, 4.3.4] Small's word experts (1980, 1983; Small and Rieger 1982) each contain a discrimination net for deciding on the word's meaning (if it is ambiguous). Thus, word experts correspond to a form of the ordered-search hypothesis in which the search is controlled not by frequency but by a net structure that, presumably, minimizes the average distance from the root to a leaf; that is, minimizes the number of questions that have to be asked in order to choose the right meaning. Could failure to control for this possibility explain the experimental results of Tanenhaus, Leiman, and Seidenberg (1979) and Seidenberg, Tanenhaus, Leiman and Bienkowski (1982)? What predictions does the discrimination net model make that would distinguish it from the competing hypotheses? Construct and run a suitable experiment.

Exercise 4.6 [4.3, 4.3.4, 5.6, 5.6.2] Using the experimental method of Swinney (1979), test for activation and decay of the senses of words that are not disambiguated by the preceding context. When (if ever) is the final choice made? Consider the effects of dominant meanings, of phrase, clause, and sentence boundaries, and of subsequent disambiguating information of various types. Look at the work of Garrod and Sanford (1985), who addressed the similar question of when an anaphor is resolved. Explain the apparent conflict of the results of Hudson and Tanenhaus (1984) with those of Granger, Holbrook, and Eiselt (1984). If Granger *et al* are correct, then is there a psychological reality to paragraphs?

Exercise 4.7 [5.1] Lexical disambiguation can be added to Montague's PTQ

formalism (1973) by "hiding" it in the translation function g. This implies, however, that g has access to the information it needs, such as certain other words in the sentence (*cf.* the concept of visibility in Polaroid Words). Can this be done by modifications to the top level of the formalism, or is a back-door method like Polaroid Words, semi-transparent to the translation level, required? In either event, provide a suitable formalization.

Exercise 4.8 [5.3.2, 8.3.3] Polaroid Words bear a superficial similarity to a simulated annealing system, or Boltzmann machine (Kirkpatrick, Gelatt and Vecchi 1983; Fahlman, Hinton and Sejnowski 1983; Smolensky 1983; Feldman 1985). What are the similarities and differences? (*Hint:* consider time.) Program a simulated annealing model of lexical disambiguation, in which all words are disambiguated at the same time. In the graph view of simulated annealing, nodes of the system will be Polaroid Words and arcs will represent visibility. Labels on the arcs will note syntactic and selectional constraints on meaning. Can the system be modified so that the nodes can be added one at a time as the sentence is processed from left to right, the system being re-annealed for each new word?[15]

Exercise 4.9 [5.3.4, 7.2.7] Develop a principled way of deciding when a PW should be required to make its final decision, even if it has to guess.

Exercise 4.10 [4.3.1, 4.3.2, 5.2.2, 5.6.1–3] [16]People's semantic associations don't always seem to be what marker passing in an elegantly organized frame system says they ought to be. In particular, people have WORD ASSOCIATIONS, which are not usually based on semantic closeness but rather on vaguer relationships such as "frequently-associated-with". These associations are generally strong and reliable, and much the same across individuals. The following examples are from Jenkins 1970:

(9-22)	priest–church	music–song	red–white
	quiet–loud	moon–stars	command–order
	doctor–nurse	eagle–bird	citizen–U.S.A.
	cheese–crackers	cabbage–head	stem–flower
	whistle–stop	working–hard	

Some of the associations, such as *doctor–nurse*, may be due to the frequency with which the two concepts named co-occur; but others are synonym or antonym pairs (*quiet–loud*, *command–order*), or phrase completions (*whistle–stop*).

We can thus distinguish two possible types of spreading activation: that between semantically close concepts, and that between associates such as those above. Two questions arise: Most of the research on semantic priming effects used stimulus

[15]For approaches to parsing in a connectionist system, see Jones 1983, Selman 1985, Selman and Hirst 1985, 1987, Pollack and Waltz 1982, Cottrell 1985b, Waltz and Pollack 1985, Jones and Driscoll 1985.

[16]I am grateful to Gary Cottrell, Rusty Bobrow, Margery Lucas, Mike Tanenhaus, and Ken Forster for discussion on the points in this exercise.

pairs many or all of which were associates rather than close concepts; is there in fact any spreading activation at all between the latter? If there are two separate kinds of spreading activation, are both of them used in lexical disambiguation?

Evidence that activation spreads between both kinds of pairs is given by the work of Fischler (1977), who found semantic priming effects both for associates such as *jump–rope* and for semantically related pairs, such as *cave–mine*, which were not associates but which shared more semantic features than control pairs such as *bread–stem*. Evidence that association links are used for disambiguation comes from semantic garden-path sentences *(section 4.3.2)*. However, it may well be that the kind of activation is different in the two cases. De Groot (1983) found no facilitation for associates of associates; that is, the associations *bull–cow* and *cow–milk* did NOT facilitate *bull–milk*. This kind of activation may thus spread a distance of exactly one link. Lupker (1984) found that semantic facilitation was much smaller than associative facilitation; Kintsch and Mross (1985) had similar results comparing associative facilitation with that from words that were thematically related but not particularly close semantically, such as *plane–gate* (**aircraft–airport-doorway**). Moreover, it seems that semantic garden paths occur just when the misleading prime is an associate to the ambiguous word, and not when there is only semantic closeness (Michael Tanenhaus, personal communication).

At present, Polaroid Words have no access to association norms. If these could be added to the Frail knowledge base, with appropriate marker passing between them, would they help or hinder disambiguation? What other effects, good or bad, could they have in the knowledge base? What would the effects be of a rule that allowed PWs to jump to a conclusion only in the case of a link by association and not one of general semantic closeness?

Experiment: Test for word association links between the negative prime and the misinterpreted word in the garden-path examples of section 4.3.2. Can you use the word association norms of Postman and Keppel (1970) to construct new semantic garden paths? Can you construct any garden-paths that don't have an associate link? What purpose (if any) does word association serve in the human mind; or of what is it an artifact? What side-effects does it have on cognitive processes?

9.5 Structural ambiguity

Exercise 5.1 [1.1.3, 6.3.2, 6.3.3] Marcus (1980: 228–234) has suggested that if semantic and syntactic biases of equal strength conflict in a sentence, the sentence is ill-formed. Test this suggestion experimentally. Since such biases vary widely from one individual to another, you will first need to devise a method for measuring the strength of an individual's syntactic and semantic biases. In your experiment, the subjects' biases are measured and then their judgments are taken on a set of sentences with bias conflicts. You will be looking to see whether a sentence is judged ill-formed whenever the same subject rated each of its component biases equal. Careful construction of the experimental materials will be crucial. You

must also take care to avoid the inherent pitfalls of a within-subjects experimental design.

Exercise 5.2 [6.2.1, 6.2.3, 7.3, 7.4.3] Punctuation can be an important cue to structural disambiguation—compare (9-23) and (9-24):

(9-23) Despite Ross's promises, to Nadia he seems as unreliable as ever.

(9-24) Despite Ross's promises to Nadia, he seems as unreliable as ever.

but Absity and the Semantic Enquiry Desk do not use it at present. Make an inventory of cases in which punctuation can be helpful or critical to disambiguation, considering in particular clause-final participles and when they do and don't require a preceding comma. Show how the SED could take account of punctuation.

Exercise 5.3 [6.2.5] I have hypothesized that a closed constituent may be reopened if no presently open constituents admit the attachment of the current constituent:

(9-25) ?Many students failed that were expected not to.

This is not a sufficient condition, however:

(9-26) *Many students failed of the negligent professor.

What are necessary and sufficient conditions for reopening a constituent? How may a Marcus parser be modified so that it can reopen constituents when necessary? What are the ramifications of reopening constituents for Absity, Polaroid Words, and the Semantic Enquiry Desk?

Discontinuous constituents are a particular problem for Absity, which is always eager to close a constituent and add it to the sentence structure, since Polaroid Words and the Semantic Enquiry Desk need well-formed partial results to do their work. This is not a good strategy for flat-structured languages such as Warlpiri and Guugu Yimidhirr, in which constituents may be fragmented arbitrarily, and one cannot know until the sentence is over whether or not more pieces of a constituent may turn up. For example, in the Warlpiri sentence (9-27) (from Nash 1980), the case filler **small child** is expressed by two words that have much of the sentence in between them; the case markings serve to tie the two words together.

(9-27)

kurdu-jarra-rlu	ka-pala		maliki	wajili-pi-nyi	wita-jarra-rlu
child-dual-erg	aux:pres-3dual(subj)		dog	chase-NPast	small-dual-erg

'The two small children are chasing the dog.'

What changes need to be made to Absity so that it can keep its partial results well-formed but allow later modifications to them?[17] Do the changes reduce the available information that PWs and the SED need? If so, what problems are created?

[17] See Brunson 1986b for a deterministic parser for Warlpiri; see also Johnson 1985 on parsing discontinuous constituents in Guugu Yimidhirr.

Do speakers of such languages avoid ambiguities that this might otherwise make difficult? Do they create ambiguities that are easy to resolve in these languages but would be difficult in English?

Exercise 5.4 [6.3.3] Is Ford, Bresnan, and Kaplan's theory of closure (1982) falsifiable?

Exercise 5.5 [7.2.2] We have said that PPs with prepositions such as *despite* can never be NP-attached and that in apparent exceptions the PP is attached to a nominalized verb with the attachment taking place before the nominalization. To parse these exceptions we will need a mechanism for detecting them. This entails having pointers from nominalizations to their verbs—from *sale* to *sell*, for example. The SED could then check for PP attachment to the nominalization by testing for its attachment to the corresponding verb. Implement a mechanism for handling nominalizations. Notice that in nominalization, pseudo-prepositions are translated to genitive constructions; thus (9-28) becomes (9-29):

(9-28) SUBJ Ross sold OBJ the book to Nadia.

(9-29) Ross's sale of the book to Nadia

Exercise 5.6 [6.3.3] Brunson (1985, 1986a) has observed that preference for verb phrase or noun phrase attachment of a locative prepositional phrase varies according as the subject of the sentence is an AGENT or an EXPERIENCER or both. (Brunson assumes Chomsky's (1982) theory of government and binding, in which an NP can play more than one THETA-ROLE—in our terms, be in more than one case slot.) This effect is especially strong with causative verbs, for which the preference for VP attachment if the subject is an AGENT overrides semantic anomaly. Examples:

(9-30) Ross baked the cake in the freezer.
*(*Ross *is the* AGENT *of* bake*; the PP is VP-attached, i.e., the baking took place in the freezer.)*

(9-31) Ross saw the man in the park.
*(*Ross *is both* AGENT *and* EXPERIENCER*; there is no preference in the attachment of the PP.)*

(9-32) Ross knew the man in the park.
*(*Ross *is the* EXPERIENCER *of* know*; the PP is NP-attached to* the man.*)*

Brunson's explanation is that verbs that take an AGENT tend to be more concrete actions that can be more easily located spatially, while verbs that take EXPERIENCER subjects tend to express more abstract relations that are harder to locate in any spatial sense. Is there any similar effect caused by the concrete–abstract continuum for nouns, and, if so, how does it interact with the effect of the agentive–experiential verb continuum (Brunson 1985, 1986a)?

Exercise 5.7 [7] We have said nothing about the effects in structural disambiguation of the length of the constituents involved. Frazier (1978) noted that VP attachment is preferred for (9-33):

(9-33) Joe returned the book for Susan.
(i.e., it was for Susan that Joe returned the book.)

but lengthening the intervening NP seems to change the preference:

(9-34) Joe returned the sunglasses, the wok, some gaudy posters, and the book for Susan.
(i.e., the book for Susan, the wok, etc, were returned by Joe.)

That is, there is a preference for local right attachment over distant minimal attachment, distance being measured in words, not constituents.

Frazier (1978; Frazier and Fodor 1978) proposed a two-stage model to account for this effect; the first stage was limited in its dealings to a short stretch of words. This model has been heavily criticized (*e.g.*, Wanner 1980; Ford, Bresnan, and Kaplan 1982). However, there seems to be something essentially right about incorporating an account of the effects of constituent length in a model of disambiguation, but neither the present work nor other critics of Frazier do this. Make a study of these effects. Can they be accounted for within the framework of the Semantic Enquiry Desk? How would its rules have to be modified? Are changes to the parser itself called for?

Exercise 5.8 [5.5, exercise 1.7] It may be regarded as an inconsistency that in section 5.5 we were insistent upon taking selectional restrictions out of the lexicon and putting them in the knowledge base, while in sections 7.2 and 7.3 we had no qualms about putting case preference information in the lexicon rather the knowledge base. We did this, of course, because we were using Ford, Bresnan, and Kaplan's results (1982). But the question obviously arises as to whether the preferences are a strictly lexical matter, or whether they actually reside in the concept underlying the lexeme.

Evidence that the latter is correct comes from the work of Kurtzman (1984). In Kurtzman's experiments, the verb of a structurally ambiguous sentence was replaced by a made-up word:

(9-35) The official shemlarked the student that the teacher had failed.

Subjects, told that the experiments concerned the teaching of foreign vocabulary, were given one of two definitions for the novel word:

(9-36) *(a)* to shemlark: to inform a person of some very upsetting news.
(b) to shemlark: to inform a very upset person of news.

Subjects who received definition *(a)*, which emphasized the news, tended to interpret (9-35) as *the student was shemlarked by the official the news that the teacher had failed*—that is, their preference was:

(9-37) [AGENT *shemlark* PATIENT MESSAGE]

By contrast, the second definition, focusing on the person, produced a preference for the other structure: *the student that the teacher failed was shemlarked*:

(9-38) [AGENT *shemlark* PATIENT]

The claim is that when parsing the sentences with novel words, subjects did not merely take the lexical preferences of a corresponding English word—the definitions were constructed so that there was no such word—but rather assigned sentence structure on the basis of what information might reasonably be provided in a sentence referring to the novel concept.

This result is intuitively appealing. It predicts that two verbs whose underlying concepts are the same, or nearly so, should have the same preferences, and, moreover, that preferences should be the same across languages. Are these predictions borne out? Can you find counterexamples? Clearly no significant change need be made to the Semantic Enquiry Desk to have it look to Frail instead of the lexicon for preferences, but what are the ramifications for Frail? Regardless of the psychological reality of conceptual preferences, is perhaps the lexicon still the best place to store case preferences in an NLU system?

Exercise 5.9 [5.6.2] When looking at lexical ambiguity, we considered at length the question of when the final choice of meaning is made. We didn't look at the same question for structural ambiguity because there the resolution point came as a given: resolution occurs exactly when the parser asks for the advice of the Semantic Enquiry Desk. However, there are many other models that could be conceived and there is much disagreement among researchers in the area. On the basis of his experiments, Kurtzman (1984) proposed a model in which competing structures are built in parallel; the point at which one is selected as the final interpretation depends on the concepts and context involved. On the other hand, Ferreira (1985) provides experimental evidence for a contrary model, in which a syntactically preferred structure is built first; if it is found implausible, or if the parser is garden-pathed, the system backs up and attempts another structure.

Ferreira's experimental material was somewhat flawed.[18] However, she makes the interesting suggestion that her results support a model in which the processor may look in the discourse model for a potential referent for a possible NP such as *the horse raced past the barn* much as the Semantic Enquiry Desk, with the Principle of Referential Success, will ask Frail if it can find a referent.

Investigate the question of at what point structural disambiguation occurs in humans. Are there speed differences that may be attributable to the need in some cases to access world knowledge to test for referential success? Could Kurtzman's conceptual factors also affect the point at which lexical disambiguation occurs?

[18] Ferreira claimed that the presuppositions of reduced relative clause sentences such as:

(i) The horse raced past the barn fell in a puddle.

were satisfied in her 'neutral context' condition; in the example she gives, they are not.

Exercise 5.10 [7.3.1, 7.3.2] Experiments by Tanenhaus, Stowe, and Carlson (1985) suggest how lexical expectations may be used by our gap-finding procedure. Subjects were given sentences with a false gap after the verb, and their reading times were measured at that point. There were four kinds of sentence tested: those in which the available filler was plausible at the false gap and those in which it was not, crossed with those with a verb whose lexical expectations were for an object to be supplied, and those in which the preference for the verb was intransitive. Examples (the false gap is marked ♦; the true gap is marked ◊):

(9-39) *Transitive expectations, plausible filler:*
I wonder which story the teacher told ♦ the children about ◊.

(9-40) *Transitive expectations, implausible filler:*
I wonder which book the teacher told ♦ the children about ◊.

(9-41) *Intransitive expectations, plausible filler:*
I wonder which patient the nurse hurried ♦ the doctor towards ◊.

(9-42) *Intransitive expectations, implausible filler:*
I wonder which bed the nurse hurried ♦ the doctor towards ◊.

When the verb expectations were transitive, reading times were longer for the implausible filler; when they were intransitive, plausibility did not affect reading time.

This suggests a model in which the parser hypothesizes a gap after the verb only if the verb is expecting a case flagged by object position. Thus the false gap was never considered in (9-41) or (9-42). In (9-39) and (9-40), it IS considered, and the attempt to fill it with something implausible, in (9-40), takes longer. Later on in these sentences, the parser discovers the true gap and, presumably, adjusts its analysis.

We have seen that because Paragram is deterministic, it has to stake everything on the SED being able to correctly accept or reject a potential gap on semantic grounds, without knowing whether or not another gap is available further to the right. Since the SED has insufficient information to always do this correctly, Paragram sometimes finds itself embarrassed. Paragram could not adopt the model suggested by the results of Tanenhaus, Stowe, and Carlson, described above, because they imply non-determinism (*cf.* exercise 6.3). However, the number of errors could perhaps be reduced if the SED, following this model, took lexical expectations into account in its gap-finding deliberations. Would this in fact be the case, or would it simply result in a different (but not smaller) set of errors? Could expectations for cases flagged by prepositions also be used by the SED for gap-finding?

9.6 None of the above

Exercise 6.1 [4.3.4] Just as Seidenberg, Tanenhaus, Leiman, and Bienkowski (1982) found that selectional restrictions are not used to inhibit access to inappropriate meanings of a word, so Frazier, Clifton, and Randall (1983) found that some

information carried on the verb that could be used in gap finding apparently isn't. Is this a trend, or what? Perhaps accessing a word or word sense does not provide immediate access to its syntactic or semantic features. How would you test this hypothesis? (See also Frazier, Clifton, and Randall's discussion (1983: 211–216), and the conflicting results of Crain and Fodor (1985) and Stowe (1984).)

Exercise 6.2 A language generation program should try to avoid producing sentences that are gratuitously ambiguous or misleading. For example, Lytinen's MOPTRANS (1984: 112–113) *(see section 8.3.1)* translated (9-43) as (9-44):

(9-43) deux jeunes filles, dont les mains avaient été blessées par suite d'une bombe

(9-44) two young women injured by a bomb in the hands

wrongly saying that the women were holding the bomb. A better translation would have been one of these:

(9-45) two young women injured in the hands by a bomb

(9-46) two young women whose hands were injured by a bomb

To avoid this sort of error, a system should check its proposed output for ambiguity and ambiguity resolvability before "uttering" it, and, ideally, it should be able to do this with the same ambiguity-handling mechanisms used for processing input. Take your favorite language generation program and add a mechanism that lets it use Polaroid Words and the Semantic Enquiry Desk in this manner.[19]

Exercise 6.3 [7.3, 7.4] Existing grammars for Marcus parsers are unable to handle several of the syntactic constructions that we looked at in chapter 7. Add the necessary grammar rules. What changes to the parsers' structure do you find necessary? Consider in particular the gap-finding problems of section 7.3.1, which, I claim, are proof that parsing without backtracking is impossible in principle.

[19] See also Taha 1983 for a few heuristics for out-of-context ambiguity avoidance.