

Interpretation Under Ambiguity

Peter Norvig
University of California, Berkeley

Introduction

This paper is concerned with the problem of semantic and pragmatic interpretation of sentences. We start with a standard strategy for interpretation, and show how problems relating to ambiguity can confound this strategy, leading us to a more complex strategy. We start with the simplest of strategies:

Strategy 1: Apply syntactic rules to the sentence to derive a parse tree, then apply semantic rules to get a translation into some logical form, and finally do a pragmatic interpretation to arrive at the final meaning.

Although this strategy completely ignores ambiguity, and is intended as a sort of strawman, it is in fact a commonly held approach. For example, it is approximately the strategy assumed by Montague grammar, where 'pragmatic interpretation' is replaced by 'model theoretic interpretation.' The problem with this strategy is that ambiguity can strike at the lexical, syntactic, semantic, or pragmatic level, introducing multiple interpretations. The obvious way to counter this problem is as follows:

Strategy 2: Apply syntactic rules to the sentence to derive a set of parse trees, then apply semantic rules to get a set of translations in some logical form, discarding any inconsistent formulae. Finally compute pragmatic interpretation scores for each possibility, to arrive at the 'best' interpretation (i.e. 'most consistent' or 'most likely' in the given context).

In this framework, the lexicon, grammar, and semantic and pragmatic interpretation rules determine a mapping between sentences and meanings. A string with exactly one interpretation is unambiguous, one with no interpretation is anomalous, and one with multiple interpretations is ambiguous. To enumerate the possible parses and logical forms of a sentence is the proper job of a linguist; to then choose from the possibilities the one "correct" or "intended" meaning of an utterance is an exercise in pragmatics or Artificial Intelligence.

One major problem with Strategy 2 is that it ignores the difference between sentences that seem truly ambiguous to the listener, and those that are only found to be ambiguous after careful analysis by the linguist. For example, each of (1-3) is technically ambiguous (*with* could signal the instrument or accompanier case, and *port* could be a harbor or the left side of a ship), but only (3) would be seen as ambiguous in a neutral context.

- (1) I saw the woman with long blond hair.
- (2) I drank a glass of port.
- (3) I saw her duck.

Lotfi Zadeh (personal communication) has suggested that ambiguity is a matter of degree. He assumes each interpretation has a likelihood score attached to it. A sentence with a large gap between the highest and second ranked interpretation has low ambiguity; one with nearly-equal ranked interpretations has high ambiguity; and in general the degree of ambiguity is inversely proportional to the sharpness of the drop-off in ranking. So, in (1) and (2) above, the degree of ambiguity is below some threshold, and thus is not noticed. In (3), on the other hand, there are two similarly ranked interpretations, and the ambiguity is perceived as such. Many researchers, from Hockett (1954) to Jackendoff (1987), have suggested that the interpretation of sentences like (3) is similar to the perception of visual illusions such as the Necker cube or the vase/faces or duck/rabbit illusion. In other words, it is possible to shift back and forth between alternate interpretations, but it is not possible to perceive both at once. This leads us to Strategy 3:

Strategy 3: Do syntactic, semantic, and pragmatic interpretation as in Strategy 2. Discard the low-ranking interpretations, according to some threshold function. If there is more than one interpretation remaining, alternate between them.

Strategy 3 treats ambiguity seriously, but it leaves at least four problems untreated. One problem is the practicality of enumerating all possible parses and interpretations. A second is how syntactic and lexical preferences can lead the reader to an unlikely interpretation. Third, we can change our mind about the meaning of a sentence—"at first I thought it meant this, but now I see it means that." Finally, our affectual reaction to ambiguity is variable. Ambiguity can go unnoticed, or be humorous, confusing, or perfectly harmonious. By 'harmonious,' I mean that several interpretations can be accepted simultaneously, as opposed to the case where one interpretation is selected. These problems will be addressed in the following sections.

Syntactic and Lexical Preferences

When we take pragmatics into account, some sentences have a prohibitively large or infinite number of interpretations. Consider the following:

(4) He seems older now.

Here *he* can refer to one of several billion males, and *now* can refer to one of an infinite number of time points. Thus, while syntax and semantics may be producing discrete lists of possibilities, it seems that pragmatic interpretation must operate by proposing likely interpretations, rather than enumerating all possibilities and then choosing among them. Hobbs (1983) has argued that enumeration should be minimized even in syntactic and semantic analysis. He considers sentences with multiple lexical ambiguities, like the following, which has at least six possible interpretations:

(5) He put the dough in the bank.

put(male,money,river-side)
put(male,money,financial-institution)

put(male,money,snow-or-cloud-mass)
put(male,flour-mixture,river-side)
put(male,flour-mixture,financial-institution)
put(male,flour-mixture,snow-or-cloud-mass)

Hobbs suggests that this list obscures the real issue: that there are two points of ambiguity (in this case, lexical ambiguity), one with two possible interpretations and one with three. He would represent the semantics of the sentence with a single logical form:

male(x) & put(x,y,z) & dough(y) & bank(z)

Pragmatic interpretation of the sentence then consists of solving for x, y, and z; the paradigm is now one of solving problems rather than making choices. The points of ambiguity are isolated, even though the solution of the problem will in general involve considering several factors at the same time. The solution will be guided by rules like the following, along with various pragmatic preference rules (this varies in detail from Hobbs' proposal):

dough(y) → money(y) | flour-mixture(y)
bank(z) → river-side(z) | financial-institution(z) | snow-or-cloud(z)

There are several complications. First, some sentences, like *I saw her duck*, are genuinely ambiguous, and would result in two distinct syntactic and logical forms, even in Hobbs' system. Second, some predications can be "pragmatically strengthened" even when they are unambiguous. Hobbs' example is that the predicate *went* in *I went to Tokyo* can be strengthened to *flew* in most contexts. By representing ambiguity with vague predicates, the same mechanism can handle ambiguity and pragmatic strengthening.

If we combine Hobbs' approach with the Necker-cube intuition, we get the following, which is similar to the approaches used by several recent authors, including Hobbs & Martin (1987), Stallard (1987), and Charniak and Goldman (1988):

Strategy 4: Do lexical, syntactic and semantic analysis to produce one (or occasionally more) neutral representation of the input, which can contain ambiguous and vague predications. Pragmatics then attempts to 'solve' for the ambiguous predications and some of the vague ones. Each solution has a ranking, and, as in Strategy 3, the low ranking solutions are discarded, and the final interpretation alternates between the high ranking one(s).

Given this approach, we still are no closer to specifying what makes us prefer one interpretation over another. It turns out that choosing the syntactically well-formed parse with the highest pragmatic likelihood is insufficient to mimic human performance. One area where the naive strategy breaks down is the so-called garden path sentence.

A garden path sentence invites the listener to consider one possible parse, and then at the end forces this parse to be abandoned in favor of another. Listeners are conscious of this switch, and often have difficulty

discovering the second interpretation. A well-known example is (6), where *raced* is initially treated as a past tense verb. This analysis fails when the verb *fell* is encountered; after some difficulty *raced* can be re-analyzed as a past participle. For most informants, there is a distinct feeling of having to re-parse the sentence; it does not feel like both parses were being built up simultaneously, and the second one was considered after the first was ruled out. If such intuitions are to be trusted (and the jury is still out), this is evidence for a problem solving approach rather than a pipelined approach which interprets all possible parses in turn.

(6) The horse raced past the barn fell.

Most informants find (7) to be much less of a garden path than (6). My analysis of this hinges on the assertion that the difficulty in processing (6) is in abandoning the parse of *raced past the barn* as a VP, and replacing it with a parse as a past-participle modifier. This is especially difficult for the reader who spends time vainly trying to come up with an interpretation of *barn fell* as an NP. In (7), there is ample evidence that *fell down and broke its leg* is a VP, which makes it easier to re-analyze *the horse raced past the barn* as a NP. Pairs like (7) and (8) are suggestive in determining how much evidence and how much time is required to force this re-analysis, but it is difficult to draw solid conclusions, since the degree to which (7) is easier to understand than (8) may be due to the additional processing time, the additional evidence for *fell* as a verb, or both.

(7) The horse raced past the barn fell down and broke its leg.

(8) The horse raced past the barn fell down.

Now consider (9-11), which are all easier to understand than (6). While (9) has the same syntactic structure as (6), it seems that *raced past the barn* is just not as good a descriptive attribute as *raced at the Kentucky Derby*. The first modifier could be true of any healthy horse, while the second describes only a top race horse. Thus, while both make good past tense verb phrases, only (9) makes a good past participle modifier. Hirst (1987) asserts that neither (10) nor (11) is a garden path. Similar forces may be at play here; it appears that *led* is a better descriptive attribute than *raced*. For whatever reasons, the reader has the ability to consider both the main verb and past participle readings for these sentences, without the feeling of a severe problem as in (6).

(9) The horse raced at the Kentucky Derby died.

(10) The horse led a long line of wagons down the road.

(11) The horse led a long way down the road fell.

Another example of this distinction is illustrated in (12) and (13). These examples are complex because *got* is highly polysemous. It can mean 'received' (as in got a raise), 'became' (got old), 'undergo' (got arrested) or 'cause/achieve' (got them arrested). In (12) the initial interpretation of the first four words is 'the boy became obese', with *fat* interpreted as an adjective. When the final word is processed, this initial interpretation has to be

abandoned.

(12) The boy got fat spattered.

(13) The boy got fat spattered on his arm.

An informal experiment in Schubert (1984) shows that (12) is a quite difficult garden path sentence, while (13) is not a garden path. One explanation for this is that having fat spattered on one's arm is the kind of experience one might be described as having undergone, while 'fat spattered' is not something one undergoes, nor is it the kind of thing one normally strives to achieve. Also, as pointed out above, (12) is difficult because it asks the reader to hypothesize *spattered* as the main verb with no additional evidence, while (13) provides the reader with the complete VP *spattered on his arm*, and also provides time to abandon the interpretation of got fat as a VP.

It seems that many ambiguous interpretations are not consciously considered because they violate selection restrictions. For example, in I drank a glass of port, the noun port is unambiguously interpreted as a kind of fortified wine, even though it has other senses. Thus, we might be tempted to formulate a principle stating that word senses which violate selection restrictions are not considered when there is another sense that satisfies the restrictions. Unfortunately, sentences (14-17), taken from Reder (1983), are evidence against this principle as it stands:

(14) The astronomer married a star.

(15) The plumber lit his pipe.

(16) The rabbi was hit on the temple.

(17) The hay farmer drank through a straw.

In each of these, the final noun has one lexical sense that satisfies the selection restrictions. However, there is another sense that is suggested first, and which stubbornly refuses to go away. Thus, as in garden path sentences, there is only one sensible interpretation, but that interpretation is difficult to arrive at. No such effect exists in sentences (18-21):

(18) The star married the astronomer.

(19) The pipe was lit by the plumber.

(20) He was hit on the temple, was the rabbi.

(21) He was drinking through a straw, was the hay farmer.

These data are often explained by assuming that all lexical senses are initially considered when a word is processed, and that a spreading-activation-like process selects for senses that are associated with concepts in context. In the presence of strong associations, as in (14-17), this process can temporarily override selection restrictions, while in (18-21) there are no strong associations at the time the ambiguous word is processed, so there is no incongruous effect.

As another example, experiments by Rayner, Carlson, and Frazier (1983) showed that (22) below was interpreted by most subjects as being silly, meaning that either the walls were painted with a crack-like pattern, or that

cracks were somehow the instrument of painting. Similarly, Hirst (1987) says that informants interpret (23) as a bizarre kind of baking in the freezer, rather than a normal oven-baking of a cake that had been in the freezer. In both cases, the informants missed the interpretation where the PP attaches to the preceding NP. Similarly, in (24) the *by* adjunct was interpreted by informants as a passive marker, not as a locative PP.

(22) The landlord painted all the walls with cracks.

(23) Ross baked the cake in the freezer.

(24) Ross was told what to do by the river.

All this evidence suggests that Strategy 4 is insufficient as it stands. In each of these examples informants discard—or delay considering—a semantically plausible interpretation, because they fixate on the 'wrong' syntactic parse or lexical sense, even though the resulting semantic interpretation is improbable. Strategy 4, unlike the informants, would have no trouble coming up with the semantically probable solution. For example, the following translations for (14) and (18) would both be handled equally easily by a Strategy 4 problem solver like the one outlined by Hobbs.

(14') $\text{star}(x) \ \& \ \text{married}(x,y) \ \& \ \text{astronomer}(y)$

(18') $\text{astronomer}(y) \ \& \ \text{married}(y,x) \ \& \ \text{star}(x)$

Thus, what Strategy 4 is missing is a way of taking timing considerations into account, and a way of stating preferences for certain parses, even when those parses do not correspond to the best interpretation. Strategies like Kimball's (1973) or Frazier and Fodor's (1978) try to account for phenomena like these in terms of general syntactic preference principles, by appealing to performance issues such as limits on available memory space. Schubert (1984, 1986) provides a good summary of these principles. Kimball's principal of Right Association says that modifiers tend to attach to the most recent phrase they could possibly modify. Thus, in (25), the preferred reading is that for *Mary* modifies *selected*, not *book* or *bought*. In (26), however, the preferred reading has the PP modifying *carried*, not *groceries*. This is explained by the principal of Minimal Attachment, which prefers parses that use the longest rewrite rules, and thus result in a parse with fewer nodes. If we assume a grammar which includes the rules $\text{VP} \rightarrow \text{V NP PP}$ and $\text{NP} \rightarrow \text{NP PP}$, then attaching the PP to the V rather than the NP minimizes the number of nodes. This analysis presupposes that Minimal Attachment takes precedence over Right Association, and that *carried* subcategorizes for the $\text{VP} \rightarrow \text{V NP PP}$ rule, while *bought* does not.

(25) John bought the book which I had selected for Mary.

(26) John carried the groceries for Mary.

Very general strategies like these are necessarily tied to particular assumptions about the underlying grammar. For example, Minimal Attachment makes no sense in categorial grammar, where all parses have the same number of nodes. But the crucial point is that these general principles just do not hold up. It is quite easy to choose different lexical items and different

situations in which the preferred interpretations reverse from those predicted by the principles. Thus (Schubert's examples again), the pairs (27-28) and (29-30) each have the same verb and preposition, but a different preferred attachment for the PP.

- (27) John met the girl that he married at a dance.
- (28) John met the girl that he saw at a dance.
- (29) John saw the bird with the powerful beak.
- (30) John saw the bird with the powerful binoculars.

At best then, syntactic preferences are only one factor that must be considered in arriving at the best interpretation, and a rather weak one at that. Schubert and especially Kurtzman (1984) come to this conclusion.

What we need is a strategy that accounts for the preferences shown above, while maintaining the other advantages of Strategy 4:

Strategy 5: Do lexical, syntactic, and semantic analysis on a word-by-word basis, identifying points of ambiguity along the way, and using all sources of evidence to rank alternatives. Evidence for a particular choice can include lexical frequency preferences, pragmatic associations, and other factors outside of the logical form. A high-ranking interpretation can be accepted (and its alternatives discarded) before the parse is complete, if its score remains sufficiently above the alternatives for a sufficient amount of time. In addition, if at any point there are more than a maximum number n ($n = 3$?) alternatives, discard the lowest ranking alternative, even if its score is close to others. At the end, alternate between the highest ranking interpretations, as before.

Mutually Compatible Interpretations and Connotations

Consider the following quote from Richard Parsons, of the American Fur Industry Inc., on their new advertising slogan *Fur is for Life*:

"It has a good sound, a good connotation. Yes, they last a long time. Yes, they're a good product. Yes, furs support wildlife conservation."

Parsons (although not a professional linguist) is making a claim about language use: that the proper or intended meaning of a phrase can be a combination of a number of interpretations and connotations. The strategies we have developed so far assume that the reader eventually arrives at a single interpretation, or a Necker alternation between interpretations. But Parsons is saying that his slogan *Fur is for Life* is different. The slogan seems to have two primary interpretations, (31) and (32) below. But it also has important connotations, listed as (33-35), as well as another interpretation, (36), that Parsons presumably wants the public to ignore.

- (31) Fur lasts a lifetime.
- (32) The fur industry is pro-conservation.
- (33) Fur wearers are lively.
- (34) The recipient of a fur may become indebted to the giver for life.
- (35) Life is a good thing; hence fur is a good thing.

- (36) Fur, while on an animal, protects its life.

Although we would not be likely to say that any of (33-35) are good candidates for the final interpretation, it seems that the intended effect of the slogan is for the reader to entertain some or all of these simultaneously. While this is a radical departure from the Hockett/Jackendoff Necker-cube theory of 'one interpretation at a time,' it appears to be quite common in poetry, politics, and advertising (see Burli-Storz, 1980).

To take another example, a pop song by David Byrne proclaims *We are creatures of love*. This can be taken as having the three interpretations listed below, but colleagues who were familiar with the song report having no feeling of switching between these interpretations.

- (37) We are born as a consequence of sexual love.
 (38) We have souls that contain or are composed largely of love.
 (39) We are possessed or controlled by the force of love.

One could argue that this is a case of vagueness, rather than ambiguity. That would mean that *of* is analyzed as simply 'related to,' rather than the more specific 'consequence of,' 'composed of,' or 'controlled by' readings. Similarly, *love* would be analyzed as the vague 'sexual, romantic, or parental love.' However, this analysis doesn't explain how the rest of the song can coherently refer to the three themes mentioned above.

The facts are admittedly slippery; I am suggesting that alternative parses can sometimes be combined into one interpretation, but it is hard to distinguish between distinct parses that have been combined together, and a vague interpretation that has several possible entailments. Also, it is notoriously hard to introspect about the phenomenology of these cases. Perhaps the following example, from the last line of Gerard Manley Hopkins' *God's Grandeur* will be more compelling:

- (40) Because the Holy Ghost over the bent
 World broods with warm breast and with ah! bright wings.

The word *broods* is lexically ambiguous between 'to sit on eggs to hatch them' and 'to think long and deeply or resentfully.' This is clearly not a case of vagueness. Yet it seems that the most natural interpretation is of a bird-like god sitting on an egg-like world (or world-like egg), pensively surveying his creation, and waiting for it to come to fruition. This interpretation clearly involves no Necker-like alternation between senses for broods; rather, it involves a simultaneous synthesis of two images.

Note that not just any images can be superimposed this way; if the world were flat, or if eggs were cubical, the combined image would not work. It is permissible to combine the images even though the world is quite a bit larger and composed of different material than the average egg, and even though the prototypical image of God does not include wings.

Poetry, like advertising, seems to sanction this superimposition of distinct parses. To support this claim, I opened a poetry anthology at random, finding the opening line of to Dylan Thomas' poem *In the Beginning*: "In the

beginning was the three-pointed star." As the rest of the poem makes clear, the three-pointed star should be taken as referring to a stellar body in primordial space, to the light in God's performative speech act "Let there be light," to the star of Bethlehem, and to the Holy Trinity. There does not seem to be a clear feeling of shifting between these referents; rather they seem to be entertained simultaneously.

Lakoff and Turner (1988) cite, but do not fully analyze, another Dylan Thomas poem, *Do not go gentle into that good night*:

- (41) Do not go gentle into that good night,
Old age should burn and rave at close of day;
Rage, rage against the dying of the light.

Understanding this passage requires knowledge of at least six metaphors for life and death. While these metaphors offer conflicting views on the nature of death, there is no feeling of having to switch between them in understanding the poem; they are all active at once. In fact, metaphors (42-47) are all used in the interpretation of the six words *go gentle into that good night*: (42) for *go*, (43) for *gentle*, (44) for *into*, (45) for *good night*, and (46) for *night*. Thus, the word *night* is being used simultaneously as a time, a destination, a container, and an adversary, all without promoting a conscious feeling of Necker-like switching.

- (42) Life is a journey.
- (43) Life is a struggle; death is an adversary.
- (44) Life is 'here'; death is a another world.
- (45) Death is sleep.
- (46) A lifetime is a day; death is night.
- (47) Life is a fire that blazes and burns out.

At this point let us try to modify Strategy 5 to account for these new findings. There are two possibilities; we can treat the combination of two interpretations as an abnormality, and try to show how it can be sanctioned, or we can treat it as the new basic interpretation mechanism, and try to show how it can be constrained.

Strategy 6a: The Conservative Simultaneity Strategy: Amend Strategy 5 to allow a simultaneous amalgam of two or more competing top-ranked interpretations, but only when sanctioned by some as-yet-unspecified factors, and only when the result is a coherent combination of the two.

Strategy 6b: The Radical Simultaneity Strategy: Always try to combine top-ranking interpretations into one image. When a coherent combination is impossible, alternate between interpretations as in Strategy 5.

To try to choose between the two, we will first consider Strategy 6b, as it is applied to sentence (48):

- (48) The chicken is ready to eat.

(48') $\text{chicken}(x) \ \& \ \text{ready}(x,e) \ \& \ \text{eating}(e) \ \& \ (\text{agent}(e,x) \ | \ \text{patient}(e,x))$

Using Strategy 6b, we could combine the two interpretations simply by accepting both parts of the disjunction, yielding 'the chicken is ready to eat the chicken.' This is by no means a normal interpretation of (48), so we have an argument against 6b. However, that argument only goes through if the proposed logical form (48') is accurate. Suppose we use the following logical form instead:

$\text{chicken}(x) \ \& \ \text{ready}(x,e) \ \& \ \text{eating}(e) \ \& \ (\text{agent}(e,x) \ \& \ \text{alive}(x) \ \& \ \text{location}(e,\text{barnyard}) \ \& \ \text{patient}(e,\text{seed}) \ \& \ \dots) \ | \ (\text{patient}(e,x) \ \& \ \neg\text{alive}(x) \ \& \ \text{location}(e,\text{table}) \ \& \ \text{agent}(e,\text{human}) \ \& \ \dots)$

Then we have two interpretations that cannot be combined coherently, neither under Strategy 6b nor 6a. Thus, we see that for 6b to be feasible, we need to insist on full script— or frame—like semantic interpretations, complete with default assumptions. We need a rich set of defaults to rule out unwanted unification of the two interpretations, even though we want to allow the possibility of overriding some of the defaults, as in "The chicken on the table is ready to eat her asparagus."

Let's try an example that does not bring as much background knowledge into play:

(49) She opened the door with a key.

The ambiguity is between with a key as an instrument of opening, and as a modifier of the door. Here there seems to be nothing to stop 6b from accepting both interpretations for the phrase, whereas we know that if this were the intended meaning, one would have to use something like the following:

(50) She opened the door with the key that was in/near/on it.

Thus, Strategy 6b as it stands is rejected. To evaluate Strategy 6a, we need to develop a better notion of sanctioning a combined interpretation, which we will address in the next section.

Puns

Consider the following advertisement for Flintstones brand Vitamins:

(51) We are Flintstones kids, 10 million strong and growing.

The coordinate *and growing* can modify either *kids* or *strong*, with the respective interpretations that the individual children are growing, or that the number of children is increasing. Most informants report recognizing both alternatives, but report an ability to fuse the two together into a single image where each individual child in an expanding group is growing. (However, no one interpreted strong as possibly modifying kids, perhaps because of the common idiomatic nature of the phrase *ten million strong*.) My analysis of this example is that the listener arrives at the two interpretations using something like Strategy 5, and in the process of trying to choose between them, realizes that both were intended interpretations, and successfully

superimposes the two images.

In short, (51) is a kind of pun. In a regular pun, the main point of the utterance is that the speaker has been clever, producing two meanings in one sentence. A secondary point is one of the meanings (and, for a very good pun, both of the meanings taken separately). But in (51) we have a special kind of pun, where the point is that both meanings are to be taken simultaneously. A similar example comes from another ad, for Michelin tires:

(52) Because you've got a lot riding on your tires.

Here the ambiguous phrase *got a lot riding on* is ambiguous between 'much depends on' and 'much is physically supported by,' with the resulting combined interpretation 'your family's safety while in the car depends on the tires.' Here again the reader must recognize the 'pun,' and the intended effect of combining the two interpretations, but here there is an added hitch: it is the combination of the two interpretations that resolves the phrase *a lot to 'your family'*; neither of the two interpretations strongly point to this interpretation singly, but together they do.

Let us compare these puns to the following example from Freud (1916):

(53) I met Baron Rothschild, and he treated me quite as his equal—quite famillionairely.

This is funny (Freud claims) because of the unexpected ease of combining familiarly with millionaire to create a new word meaning 'as familiarly as is possible for a millionaire.' (In German, familiär + Millionär = familionär.) Freud also presents the standard definition of joking as the ability to find hidden similarities between dissimilar things. This is amended to allow for the discovery of differences, or just "to bind into a unity, with surprising rapidity, several ideas which are in fact alien to one another." In other words, the combination of disparate ambiguous interpretations is an unusual event, but one that we have an automatic capacity for.

A remaining problem is to explain why some such ambiguities are funny, while others are not. Why is it that, to my ears at least, *the rabbi was hit on the temple* is funny, while *the plumber lit his pipe* is merely confusing? Freud claims that the laughter response is illicit by the release of suppressed violent or sexual thoughts. That explains, perhaps, why the following is a fairly good joke, while other lexical and structural ambiguities in this paper are not:

(54) She criticized his apartment, so he knocked her flat.

Minsky (1980) recasts Freud's notions into the terminology of mental agents acting as censors to violent or sexual thoughts. In Minsky's terms, certain mental agents are good at combining ambiguous interpretations, but other agents notice that this is not the normal mode of operation, and act to censor them. The laughter response serves to 'shake up' the mind, get it back on track, and post a warning to avoid such thoughts. Presumably, the simultaneous combinations that sneak by uncensored are ones that do not represent 'dangerous' modes of thought.

Simultaneous Interpretation in 'Normal' Language

There are cases of combined simultaneous interpretation which don't involve poetic license or puns. Consider the use of *book* in (55). *Book* is polysemous between a physical object, a string of words, and an abstract plot or sequence of situations. The use of *beautifully bound* refers to the physical object, *one new idea* refers to the abstract content, and *50,000 words* refers to a particular (abstract) instantiation of the content. (If the book were reprinted in paperback it would still have the same number of words, whereas if it were translated into another language, it would have a different number of words, but the same number of ideas.) All three polysemous interpretations of *book* are used simultaneously.

It can't be that *book* is a single sense referring to all these aspects, because in (56), *book* must refer only to the 'plot or sequence of ideas' sense. One could not felicitously use (56) to describe someone who had written a single book which has had a hundred copies printed, or a single book which was translated by others into a hundred languages.

(55) This book, although beautifully bound, contains only one new idea in 50,000 words.

(56) He is the author of over 100 books.

Len Talmy (1977) provides a good example of image combination in non-ambiguous language. (57) forces the reader to combine the image of a woman walking through a party with the image of a leaf wafting through the air (or something similar) to arrive at an interpretation. Talmy explains just what properties of the verb are maintained, and which are taken from the complements.

(57) She wafted through the party.

Image combination is more obvious in the case of metaphors and cliches where the derived meaning is removed than the surface form. Compare (58), which is a consistent use of metaphor, with (59). Sentence (59) provides a topological clash that cannot easily be resolved into a single interpretation, even though the meaning of the two cliches is consistent. (57) is an example of a mixed metaphor with varying effectiveness; some find it to be fine, while others report that the topology is all wrong: one should be striving to get out of a rat-hole, not out from under it.

(58) I've always been 100% behind my husband, pushing him on as best I can.

(59) I've always been at my husband's side, 100% behind him.

(60) They can't afford to get out from under the rat-hole of rent payments.

Conclusion

In this paper I have investigated several strategies for pragmatic interpretation, and have presented a new strategy which (1) accounts for the little-mentioned phenomenon of a simultaneous combination of ambiguous interpretations, (2) is not inconsistent with experimental preference results, and (3) uses a combination mechanism that is needed for non-ambiguous language as well.

Acknowledgement

The author has been sponsored by the Defense Advanced Research Projects Agency (DoD), Arpa Order No. 4871, monitored by Space and Naval Warfare Systems Command under Contract N00039-84-C-0089.

Bibliography

- Burli-Storz, Claudia. 1980. *Deliberate Ambiguity in Advertising*, Bern.
- Charniak, Eugene and Robert Goldman. 1988. "A Logic for Semantic Interpretation," *Proceedings of the Association for Computational Linguistics, 26th Annual Meeting*, Buffalo, NY.
- Frazier, Lyn and Janet Dean Fodor. 1978. "The sausage machine: A new two-stage parsing model," *Cognition*, 6(4), 291-325.
- Freud, Sigmund. 1916. *Jokes and Their Relation to the Unconscious*, Translated by James Strachey, 1960, Norton, NY.
- Hirst, Graeme. 1987. *Semantic Interpretation and the Resolution of Ambiguity*, Cambridge Univ. Press, Cambridge.
- Hobbs, Jerry. 1983. "Representing ambiguity," *Proceedings of the Second West Coast Conference on Formal Linguistics, Stanford University*.
- Hobbs, Jerry, and Paul Martin. 1987. "Local pragmatics," *Proceedings of the International Joint Conference on Artificial Intelligence*.
- Hockett, Charles F. 1954. "Two models of grammatical description," *Word*, 386-399.
- Jackendoff, Ray. 1987. *Consciousness and the Computational Mind*, MIT Press.
- Kimball, John. 1973. "Seven principles of surface parsing in natural language," *Cognition*, 2(1), 15-47.
- Kurtzman, Howard S. 1984. "Ambiguity resolution in the human syntactic parser: An experimental study," *Proceedings of Coling84*, Stanford, CA.
- Lakoff, George, and Mark Turner. 1988. *More Than Cool Reason*, to appear.
- Lyons, John. 1977. *Semantics*, Cambridge Univ. Press, Cambridge.
- Minsky, Marvin. 1980. "Jokes and the logic of the cognitive unconscious," AI memo No. 603, MIT AI Lab.
- Rayner, Keith, Marcia Carlson and Lyn Frazier. 1983. "The interaction of syntax and semantics during sentence processing: Eye movements in the analysis of semantically biased sentences." *Journal of verbal learning and verbal behavior*, 22(3), 358-374.
- Reder, Lynne M. 1983. "What kind of pitcher can a catcher fill? Effects of priming on sentence comprehension." *Journal of verbal learning and*

verbal behavior, 22.

- Schubert, Lenhart K. 1984. "On parsing preferences," *Proceedings of Coling84*, Stanford, CA.
- Schubert, Lenhart K. 1986. "Are there preference trade-offs in Attachment Decisions?," *Proceedings of AAAI*, 601-605.
- Stallard, David. 1987. "The logical analysis of lexical ambiguity," *25th Annual Meeting of the ACL*, 179-185.
- Talmy, Leonard. 1977. "Rubber-sheet cognition in language," *Chicago Linguistic Society*, 13, 612-628.
- Zwicky, Arnold M., and Jerrold M. Sadock. 1975. "Ambiguity tests and how to fail them," *Syntax and Semantics*, 4, John P. Kimball, ed., Academic Press.